

Gnome Planet - Latest News

- [FOSSGIS Conference 2025](#) (2025/04/05 07:45)

Last week I attended this year's FOSSGIS Konferenz in Münster, Germany, focusing on public transport and indoor navigation topics. Photo by FOSSGIS e.V./Sergey Mukhametov, CC-BY-SA Indoor Navigation We use indoor map rendering and indoor routing in KDE Itinerary for e.g. train stations. There were a number of interesting talks and discussions in that area: A 2.5D visualization of buildings and especially of the vertical connectivity (stairs, elevators, etc) between floors, by a team from TU Dresden, with a free implementation available. Compared to our current 2D view this makes it easier to understand how to move between floors, but it brings new challenges as well (video). Indoor localization and routing including an AR demo by the German Federal Agency for Cartography and Geodesy (BKG). The routing is done via Unity, which uses the same library underneath that we use as well. Localization is using a mix of sensors available in normal smart phones, which is particularly interesting for us. Unfortunately like most similar systems the implementation isn't published (talk). Work on importing building geometries for 400+ buildings of TU München, using IFC BIM data. That's the best way to get high-quality building geometries into OSM I think, but while there has been previous work on building an import toolchain for this, getting the actual BIM data has been elusive due to various concerns by building owners. Modelling and tagging of infrastructure in and around train stations in OSM (session notes). And of course there was also the yearly offline FOSSGIS edition of the OSM indoor meetup for connecting everyone working on that subject (translated session notes). Open Transport The other big topic for me was public transport routing, which we do in Transitous, for use in Itinerary and many other apps. There were several talks and poster sessions in that area: Katharina presented her work on an explorative long distance journey planning UI (English version from FOSDEM here). I'd really like to see Transitous become a platform to build such innovative UIs on top, so you can focus on that rather than redoing the entire data handling and routing. My own talk introducing Transitous. Felix spoke about the routing capabilities of MOTIS, the routing engine Transitou is built upon. René showed work on developing and standardizing data formats for describing public transport vehicles in great detail. That's badly needed for considering accessibility during routing, the current modelling in GTFS (basically "wheelchair=yes|no") is way to simplistic for that. A talk about DB Regio's vehicle position processing. Especially deriving delay estimates from position data is highly relevant for Transitous, but unfortunately neither their data nor their code seems to be openly available. There also has been a lot of hallway track discussion on this: Together with several MOTIS users and contributors we reviewed and updated the MOTIS feature wish list. There are ideas for collaborating on a public transport optimized geocoder, between Nominatim, Transitous and other people needing this. Currently this is usually based on OSM data, but additionally using schedule data would allow things like considering the number of lines or trips at a stop for ranking. "Meta-stations" were another subject, ie. if you select "Paris" as the destination any of its major railway stations and airports would be expected destination, rather than just the first stop inside the city boundary. GTFS data quality was a prevalent topic, including during the offline FOSSGIS edition of the Open Transport Meetup, and now also with the German Federal Agency for Cartography and Geodesy (BKG) joining everyone else in exploring the wonders of this. More eyes (and more pressure) can only help here. Possible datasets for elevation models usable for routing, for Transitous something particularly interesting for first/last mile bike routing. Plans for a Transitous sprint/hack weekend are getting a bit more concrete, with two promising options in June/July. Ideas for improving Transitous' outreach activities and social media presence/PR. Also, if you

are interested or involved in FOSS or Open Data topics around mobility/transit keep October 17/18 free, there's an exciting announcement coming up shortly. OSM The FOSSGIS conference is of course much larger than just those two topics, and while those kept me occupied most of the time I managed to pick up a few more things as well: Pirmin's talk on showing realtime data layers with MapLibre should come in handy for getting live updates onto the map view for monitoring our emergency and weather alert aggregation server. This talk also featured a Transitous live demo, something I didn't dare doing myself. Frederik's work on PostPass could turn out to be a vastly more efficient way to extract region or timezone boundary polygons than our current approach that takes hours and needs a full local OSM planet file. And as always when hanging out with people from different communities there's experiences to share on things like mentorship and funding programs as well as nowadays sadly also on defense measures to protect our infrastructure from overly aggressive AI crawler bots.

- [This Week in Plasma: polish and stability](#) (2025/04/05 02:00)

Welcome to a new issue of "This Week in Plasma"! Every week we cover the highlights of what's happening in the world of KDE Plasma and its associated apps like Discover, System Monitor, and more. This was a week of smaller improvements: lots of stability enhancements user interface upgrades. All good preparation for Plasma 6.4, which will be released in a little over two months! Notable UI Improvements Plasma 6.4.0 Sticky Notes widgets that live in a panel can now have their pop-up pinned open, just like most other panel pop-ups. (Christoph Wolk, [link](#)) Notifications no longer switch to scientific notation when showing extremely large numbers; now they always show normal numbers. (Nate Graham, [link](#)) Did a round of UI polish on the crash reporting wizard to improve its layout, appearance, and text clarity. (Thomas Duckworth, [link](#)) If you change the cursor blink rate (currently a hidden setting, but we may expose it in the GUI), this preference will now be synced to GTK-based apps too. (Kai Uwe Broulik, [link](#)) KWin's Zoom effect and Plasma's desktop mouse wheel actions are now much easier to trigger and end by scrolling with a touchpad or a mouse with a high resolution scroll wheel. (Xaver Hugl, [link 1](#) and [link 2](#)) Persistent notifications can now be sent to the notification history in case you don't want to look at them anymore, but keep them around for later, e.g. for timers. (Kai Uwe Broulik, [link](#)) Improved the visuals of the portal-based account details request dialog. (Joshua Goins, [link](#)) Sound themes can now be applied with a double-click, same as other items on System Settings' grid-based theme chooser pages. (Kai Uwe Broulik, [link](#)) Frameworks 6.14 Telegram's System Tray icons are once again Breeze-themed when using the Breeze Icon theme and Telegram version 5.12.4 or later; we had to create some symlinks to react to Telegram changing the icon names they use. (Rocket Aaron and Nate Graham, [link](#)) Notable Bug Fixes Plasma 6.3.4 Fixed a variety of small layout bugs in notifications: now URLs and long words wrap properly; the orange line on critical notifications no longer squares off the normally-rounded bottom-left corner; and internal paddings and spacings have been restored to exactly what they were in Plasma 6.2 and earlier, undoing some small layout changes unintentionally introduced with the big code refactor in Plasma 6.3. (Nate Graham, [link 1](#), [link 2](#), [link 3](#), [link 4](#), and [link 5](#)) The "Always open with this app" option in portal-based app chooser dialogs now actually works. (Nicolas Fella, [link](#)) Plasma 6.3.5 Fixed a somewhat common Plasma crash related to power-cycling screens in multi-screen setups. (Marco Martin, [link](#)) Fixed a bunch of bugs related to notifications not moving to new positions when they should under various circumstances. (Marco Martin, [link](#)) When using multiple screens, reverting a change to one of the screen's settings no longer makes the UI inappropriately show the settings of the other screen. (Oliver Beard, [link](#)) Plasma 6.4.0 Fixed a case where KWin could crash with misbehaving 3rd-party scripts. (David Redondo, [link](#)) Fixed auto-update in Discover again, as part of a larger package of work to make the notifier tray widget behave more reliably and predictably. (Aleix Pol, [link](#)) Switching pages in Plasma widgets' configuration dialogs no longer inappropriately prompts you to save or discard your changes when there were no changes! (Christoph Wolk, [link](#)) Frameworks 6.13 Fixed a bug in Kirigami that resulted in Plasma widgets' config dialogs and "Get New

[thing]" dialogs sometimes missing their header text or buttons. (Marco Martin, Link) Other bug information of note: 1 very high priority Plasma bug (up from 0 last week). Current list of bugs 18 15-minute Plasma bugs (same as last week). Current list of bugs Notable in Performance & Technical Plasma 6.4.0 System Monitor and its widgets are now capable of getting statistics from Intel GPUs. (David Redondo, link) Improved KWin's startup speed a little bit by not having it pointlessly compute font metrics for something that didn't really need to adjust with the font size anyway. (Aleix Pol Gonzalez, link) Refactored the portal-based dialogs in a way that eliminates an entire class of crashes, both actual and potential. (David Redondo, link) Made Plasma practically silent in its log output. This project is almost finished! (Christoph Wolk, link 1, link 2, link 3, link 4, link 5, link 6, and link 7) 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. You can help KDE by becoming an active community member and getting involved somehow. 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: Triage and confirm bug reports, maybe even identify their root cause Contribute designs for wallpapers, icons, and app interfaces Design and maintain websites Translate user interface text items into your own language Promote KDE in your local community ...And a ton more things! You can also help us by making a donation! Any monetary contribution — however small — will help us 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, feel free to push a commit to the relevant merge request on invent.kde.org.

- [Web Review, Week 2025-14](#) (2025/04/04 10:28)

Let's go for my web review for the week 2025-14. How crawlers impact the operations of the Wikimedia projects Tags: tech, ai, machine-learning, gpt, criticism, knowledge Unsurprisingly, Wikimedia is also badly impacted by the LLM crawlers... That puts access to curated knowledge at risk if the trend continues. <https://diff.wikimedia.org/2025/04/01/how-crawlers-impact-the-operations-of-the-wikimedia-projects/> Beyond Public Access in LLM Pre-Training Data: Non-public book content in OpenAI's Models - Social Science Research Council (SSRC) Tags: tech, ai, machine-learning, gpt, copyright, ethics We just can't leave the topic of how the big model makers are building their training corpus unaddressed. This is both an ethics and economics problem. The creators of the content used to train such large models should be compensated in a way. Between this, the crawlers they use and the ecological footprint of the data centers, there are so many negative externalities to those systems that law makers should have cease the topic a while ago. The paradox is that if nothing is done about it, the reckless behavior of the model makers will end up hurting them as well. <https://www.ssrc.org/publications/beyond-public-access-in-llm-pre-training-data-non-public-book-content-in-openais-models/> AI ambivalence Tags: tech, ai, machine-learning, gpt, copilot, criticism I somehow recognise myself in this piece. Not completely though, I disagree with some of the points... but we share some baggage so I recognize another fellow.

<https://nolanlawson.com/2025/04/02/ai-ambivalence/> Why I stopped using AI code editors Tags: tech, ai, machine-learning, gpt, copilot, learning, knowledge Even if you use LLMs, make sure you don't depend on them in your workflows. Friction can indeed have value. Also if you're a junior you should probably seldom use them, build your skill and knowledge first... otherwise you'll forever be a beginner and that will bite you hard.

<https://lucianonooijen.com/blog/why-i-stopped-using-ai-code-editors/> Pixelfed leaks private posts from other Fediverse instances - fiona fokus Tags: tech, social-media, fediverse Clearly the security practice around Pixelfed bears questioning. I'm also a bit surprise at the lack of protection of private messages in the ActivityPub protocol (even though it's a hard admittedly a hard problem).

<https://fokus.cool/2025/03/25/pixelfed-vulnerability.html> How to report a security issue in an open source project - Jacob Kaplan-Moss Tags: tech, foss, security This is considered standard practice at this point. The article does a good job explaining it and the reasoning behind it.

<https://jacobian.org/2025/mar/27/reporting-security-issues-in-oss/> The Surprise of Multiple Dependency Graphs - ACM Queue Tags: tech, dependencies, supply-chain Dependency resolution is harder than people generally expect. This is a difficult problem and is very sensitive to the context. <https://queue.acm.org/detail.cfm?ref=rss&id=3723000> A Retrospective on the Source Code Control System Tags: tech, version-control, history This paper is a look back at SCCS. This is nice to see how much progress was made in version control systems since then, it's also interesting to see how the design choices changed. <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=10821013> git-revise Tags: tech, version-control, git, tools Looks like a nice alternative to git rebase to manage patchsets. Definitely interesting if you're using something like Gerrit. With other forges... It'll depend how your reviews are handled I think. <https://mystor.github.io/git-revise.html> Gerrit, GitButler, and Jujutsu projects collaborating on change-id commit footer Tags: tech, version-control, git, codereview Could be interesting if it gets standardized. Maybe other forges than Gerrit will start leveraging the concept, this would improve the review experience greatly on those.

<https://lore.kernel.org/git/CAESOdVAspxUJKGAA58i0tvks4ZOfoGf1Aa5gPr0FXzdcywqUUw@mail.gmail.com/T/#u> Empowering WebAssembly with Thin Kernel Interfaces Tags: tech, webassembly, virtualization, portability, research This is interesting research. It shows nice prospects for WebAssembly future as a virtualization and portability technology. I don't think we'll see all of the claims in the discussion section realized though. <https://dl.acm.org/doi/pdf/10.1145/3689031.3717470> A Study of Undefined Behavior Across Foreign Function Boundaries in Rust Libraries Tags: tech, rust, safety Rust itself might bring interesting properties in term of safety. As soon as it needs to interact with other languages though the chances of undefined behavior increase drastically. This definitely pushes towards using more dynamic analysis tools to catch those. <https://arxiv.org/abs/2404.11671> Introducing Stringleton Tags: tech, rust, memory, safety Nice feature, but more interesting in its explanation is the topic of static initializers in Rust. They're clearly not a settled area in the language, that's in part because of how low level static analyzers are. <https://simonask.github.io/introducing-stringleton/> A pattern for obtaining a single value while holding a lock Tags: tech, multithreading, c++ This is indeed a nice pattern to obtain a value, brings neat advantages. <https://devblogs.microsoft.com/oldnewthing/20250127-00/?p=110809> Deadlock-free Mutexes and Directed Acyclic Graphs Tags: tech, multithreading, rust Interesting trick to check at runtime that you always acquire mutexes in the same order. <https://bertptrs.nl/2022/06/23/deadlock-free-mutexes-and-directed-acyclic-graphs.html> Anime.js | JavaScript Animation Engine Tags: tech, web, frontend, animation Interesting JS library for animation on the Web. It's nice that it seems really small. <https://animejs.com/> Learn CSS Layout The Pedantic Way Tags: tech, web, frontend, css Looks like a nice resource to deep dive into CSS layouts and really understand their behaviours. <https://book.mixu.net/css/> CSS System colors Tags: tech, web, frontend, css, colors Nice way to have a web frontend which respects the system color choices of the user. <https://anto.pt/articles/css-system-colors> Minimal CSS-only blurry image placeholders Tags: tech, web, frontend, colors This is a very smart way to create pure CSS placeholders. <https://leanrada.com/notes/css-only-lqip/> The Fifth Kind of Optimisation Tags: tech, multithreading, optimization, rust A good look back at parallelisation and multithreading as a mean to optimise. This is definitely a hard problem, and indeed got a bit easier with recent languages like Rust. https://tratt.net/laurie/blog/2025/the_fifth_kind_of_optimisation.html Ports and fat adapters Tags: tech, architecture, complexity A good reminder of why you often don't want to follow an architecture pattern to the letter. They should be considered like guidelines and depending on your technical context you should properly balance the costs. Here is an example with the Ports and Adapters pattern in the context of an ASP.NET application. <https://blog.ploeh.dk/2025/04/01/ports-and-fat-adapters/> Thoughts on ECS | Voxagon Blog Tags: tech, architecture, simulation, game Nice post about pros and cons of ECS architectures. <https://blog.voxagon.se/2025/03/28/thoughts-on-ecs.html> The manager I hated and the lesson he taught me Tags: tech, programming, engineering, leadership For sure the aforementioned manager need to fix his communication

style. That being said the core advice was indeed good. <https://www.blog4ems.com/p/the-manager-i-hated> Bye for now!

- [Akademy 2025 Call for Proposals is Now Open](#) (2025/04/03 10:04)

Akademy 2025 will be a hybrid event held simultaneously in Berlin, Germany, and online. The Call for Participation is open! Send us your talk ideas and abstracts. Why talk at #Akademy2025 Akademy attracts artists, designers, developers, translators, users, writers, companies, public institutions and many other KDE friends and contributors. We celebrate the achievements and help determine the direction for the next year. We all meet together to discuss and plan the future of the Community and the technology we build. You will meet people who are receptive to your ideas and can help you with their skills and experience. You will get an opportunity to present your application, share ideas and best practices, or gain new contributors. These sessions offer the opportunity to gain support and make your plans for your project become a reality. How to get started Do not worry about details or slides right now. Just think of an idea and submit some basic details about your talk. You can edit your abstract after the initial submission. All topics relevant to the KDE Community are welcome. Here are a few ideas to get you started on your proposal: How KDE can empower building robust communities in changing political climates Work towards KDE's goals: Streamlined Application Development Experience, We care about your Input, and KDE Needs You Giving people more digital freedom, sovereignty, and autonomy with KDE software Guides on how to participate for new users, intermediates and experts, New developments/plans for KDE Frameworks, Plasma, Applications and other projects Anything else that might interest the audience. To get an idea of talks that were accepted, check out the program from previous years: 2024, 2023, 2022, 2021 and 2020. For more details and information, visit our Call for Participation.

- [Model/View Drag and Drop in Qt - Part 3](#) (2025/04/03 08:43)

Model/View Drag and Drop in Qt - Part 3 In this third blog post of the Model/View Drag and Drop series (part 1 and part 2), the idea is to implement dropping onto items, rather than in between items. QListWidget and QTableWidgetItem have out of the box support for replacing the value of existing items when doing that, but there aren't many use cases for that. What is much more common is to associate a custom semantic to such a drop. For instance, the examples detailed below show email folders and their contents, and dropping an email onto another folder will move (or copy) the email into that folder. Step 1 Initial state, the email is in the inbox Step 2 Dragging the email onto the Customers folder Step 3 Dropping the email Step 4 The email is now in the customers folder With Model/View separation Example code can be found here for flat models and here for tree models. Setting up the view on the drag side Call `view->setDragDropMode(QAbstractItemView::DragOnly)` unless of course the same view should also support drops. In our example, only emails can be dragged, and only folders allow drops, so the drag and drop sides are distinct. Call `view->setDragDropOverwriteMode(...)` true if moving should clear cells, false if moving should remove rows. Note that the default is true for QTableView and false for QListView and QTreeView. In our example, we want to remove emails that have been moved elsewhere, so false is correct. Call `view->setDefaultDropAction(Qt::MoveAction)` so that the drag defaults to a move and not a copy, adjust as needed Setting up the model on the drag side To implement dragging items out of a model, you need to implement the following -- this is very similar to the section of the same name in the previous blog post, obviously:

```
class EmailsModel : public QAbstractTableModel {
    ~~~ Qt::ItemFlags flags(const QModelIndex &index) const override { if (!index.isValid()) return {}; return Qt::ItemIsEnabled | Qt::ItemIsSelectable | Qt::ItemIsDragEnabled; } // the default is "copy only", change it Qt::DropActions supportedDragActions() const override { return Qt::MoveAction | Qt::CopyAction; }
    QMimeData *mimeData(const QModelIndexList &indexes) const override;
    bool removeRows(int position, int rows, const QModelIndex &parent) override;
     Reimplement flags() to add Qt::ItemIsDragEnabled in the case of a valid index
     Reimplement supportedDragActions() to return Qt::MoveAction | Qt::CopyAction or whichever you want to support (the default is CopyAction only).
     Reimplement mimeData() to serialize the
```

complete data for the dragged items. If the views are always in the same process, you can get away with serializing only node pointers (if you have that) and application PID (to refuse dropping onto another process). See the previous part of this blog series for more details.☐

Reimplement `removeRows()`, it will be called after a successful drop with `MoveAction`. An example implementation looks like this: `bool EmailsModel::removeRows(int position, int rows, const QModelIndex &parent) { beginRemoveRows(parent, position, position + rows - 1); for (int row = 0; row < rows; ++row) { m_emailFolder->emails.removeAt(position); } endRemoveRows(); return true; }` Setting up the view on the drop side☐ Call `view->setDragDropMode(QAbstractItemView::DropOnly)` unless of course it supports dragging too. In our example, we can drop onto email folders but we cannot reorganize the folders, so `DropOnly` is correct. Setting up the model on the drop side To implement dropping items into a model's existing items, you need to do the following: `class FoldersModel : public QAbstractTableModel { ~~~ Qt::ItemFlags flags(const QModelIndex &index) const override { CHECK_flags(index); if (!index.isValid()) return {}; // do not allow dropping between items if (index.column() > 0) return Qt::ItemsEnabled | Qt::ItemsSelectable; // don't drop on other columns return Qt::ItemsEnabled | Qt::ItemsSelectable | Qt::ItemsDropEnabled; } // the default is "copy only", change it Qt::DropActions supportedDropActions() const override { return Qt::MoveAction | Qt::CopyAction; } QStringList mimeTypes() const override { return {QString::fromLatin1(s_emailsMimeType)}; } bool dropMimeData(const QMimeData *mimeData, Qt::DropAction action, int row, int column, const QModelIndex &parent) override; }; ☐`

Reimplement `flags()` For a valid index (and only in that case), add `Qt::ItemsDropEnabled`. As you can see, you can also restrict drops to column 0, which can be more sensible when using `QTreeView` (the user should drop onto the folder name, not onto the folder size).☐ Reimplement `supportedDropActions()` to return `Qt::MoveAction | Qt::CopyAction` or whichever you want to support (the default is `CopyAction` only).☐ Reimplement `mimeTypes()` - the list should include the MIME type used by the drag model.☐ Reimplement `dropMimeData()` to deserialize the data and handle the drop. This could mean calling `setData()` to replace item contents, or anything else that should happen on a drop: in the email example, this is where we copy or move the email into the destination folder. Once you're done, return `true`, so that the drag side then deletes the dragged rows by calling `removeRows()` on its model. `bool FoldersModel::dropMimeData(const QMimeData *mimeData, Qt::DropAction action, int row, int column, const QModelIndex &parent) { ~~~ // safety checks, see full example code EmailFolder *destFolder = folderForIndex(parent); const QByteArray encodedData = mimeData->data(s_emailsMimeType); QDataStream stream(encodedData); ~~~ // code to detect and reject dropping onto the folder currently holding those emails while (!stream.atEnd()) { QString email; stream >> email; destFolder->emails.append(email); } emit dataChanged(parent, parent); // update count return true; // let the view handle deletion on the source side by calling removeRows there } Using item widgets Example code: QListWidget QTableWidget QTreeWidget On the "drag" side☐ Call widget->setDragDropMode(QAbstractItemView::DragOnly) or DragDrop if it should support both☐ Call widget->setDefaultDropAction(Qt::MoveAction) so that the drag defaults to a move and not a copy, adjust as needed☐ Reimplement Widget::mimeData() to serialize the complete data for the dragged items. If the views are always in the same process, you can get away with serializing only item pointers and application PID (to refuse dropping onto another process). In our email folders example we also serialize the pointer to the source folder (where the emails come from) so that we can detect dropping onto the same folder (which should do nothing). To serialize pointers in QDataStream, cast them to quintptr, see the example code for details. On the "drop" side☐ Call widget->setDragDropMode(QAbstractItemView::DropOnly) or DragDrop if it should support both☐ Call widget->setDragDropOverwriteMode(true) for a minor improvement: no forbidden cursor when moving the drag between folders. Instead Qt only computes drop positions which are onto items, as we want here.☐ Reimplement Widget::mimeTypes() and return the same name as the one used on the drag side's mimeData☐`

Reimplement `Widget::dropMimeData()` (note that the signature is different between `QListWidget`, `QTableWidget` and `QTreeView`) This is where you deserialize the data and handle the drop. In the email example, this is where we copy or move the email into the destination folder. Make sure to do all of the following: any necessary behind the scenes work (in our case, moving the actual email) updating the UI (creating or deleting items as needed) This is a case where proper model/view separation is actually much simpler. Improvements to Qt while writing and testing these code examples, I improved the following things in Qt, in addition to those listed in the previous blog posts: [QTBUG-2553 QTreeView with `setAutoExpandDelay\(\)`](#) collapses items while dragging over it, fixed in Qt 6.8.1 Conclusion I hope you enjoyed this blog post series and learned a few things. The post [Model/View Drag and Drop in Qt - Part 3](#) appeared first on KDAB.

- [If your notifications look kind of stupid in Plasma 6.3.4, it's my fault](#) (2025/04/02 21:51)

This is for everyone upgrading to Plasma 6.3.4, which was released yesterday. I suspect that some of you will notice something slightly wrong with notifications; the top padding is off, causing text to look not vertically centered most of the time. This is my fault. The recent bug-fixes I made to notification spacings and paddings were backported to Plasma 6.3.4, but ended up missing a part that positions the text labels nicely when there's body text or an icon, and didn't notice this until after 6.3.5 was released. The fix was just merged and backported for Plasma 6.3.5, so unless your distro backports the fix (I've already emailed the appropriate mailing list about this) you'll have to live with slightly ugly label positioning until then. Sorry folks! My bad. Once you have the fix — either because your distro backports it or because you've waited until Plasma 6.3.5 — notification text positioning should look better again:

- [Konsole Layout Automation \(part 2\)](#) (2025/04/02 00:00)

In [Konsole Layout Automation \(part 1\)](#), I wrote about how to automate opening Konsole with different tabs that ran different commands. In this post, I'll talk about doing this for layouts. Inspiration In the past, I needed to open two connections to the same host over ssh, and change to two different directories. I opened Konsole with a layout that had two panes. Then, Quick Commands allowed me to run a command in each pane to ssh and change to the right directory. This post will outline how to achieve that and more! Goal: Launch a Konsole window with one tab that has multiple panes which run commands `□` Note For more detailed instructions on using Konsole, please see the output of `konsole --help` and take a look at the [Command-line Options](#) section of the [Konsole Handbook](#) A layout can save and load a set of panes. Unfortunately, it can't do anything else. We can, however, use profiles and the Quick Commands plugin to make the panes more useful. Use case: See the output of different commands in the same window. For instance, you could be running `htop` in one pane and open your favorite editor in another. Here's an overview of the steps: Set up a layout Use QuickCommands to run things in the panes Set up a layout Unfortunately, the online documentation for Konsole command line options doesn't say much about how to create a layout, or the format of its JSON file. It only mentions the command line flag `--layout`. Also make a note of `-e` which allows you to execute a command. Fortunately, creating the layout is pretty easy. Note that a layout is limited to one tab. It will only save the window splits, nothing else. No profiles, directories, etc. Set up a tab in Konsole with the splits you want it to have Use View -> Save Tab Layout to save it to a .json file. (I personally recommend keeping these in a specific directory so they're easy to find later, and for scripting. I use `~/konsole-layouts`). You can then use `konsole --layout ~/layout-name.json` to load konsole with a tab that has the splits you saved. Use Quick Commands to do useful things in your layout As mentioned above, you can only save splits. you can't associate a profile, or run a command directly like you can with the `tilix` or `kitty` terminals. This has been requested. In the meantime, an easy thing you can do is load a layout and then load a profile manually in each pane. This is where Quick Commands come in. These are under [Plugins - Quick Commands](#). (If you don't see this, contact your distro / the place you installed Konsole from). You can use Quick Commands to run a command in

each pane. You can also launch a profile (with different colors etc) that runs a command (part 1 showed how these might be used). Note, however, that running konsole itself from here will launch a new Konsole window. End each command with `|| return` so that you get to a prompt if the command fails. Examples `htop || return` So, after you've launched Konsole with your layout as described above, you can do this: Go to Plugins -> Show Quick Commands Add commands you'd like to run in this session. Now, focus the pane and run a command. Using these steps, I can run `htop` in one pane and `nvtop` in the other. After you've gotten familiar with tabs and layouts, you can make a decently powerful Konsole session. Combine these with a shell function, and you can invoke that session very easily. This is still too manual! You're right. This post is about automating Konsole and having to click on things is not exactly that. You can use `dbus` commands in a script to load your tab layout and then run commands in each pane without using Quick Commands. As we saw in the last post, you can use profiles to customize color schemes and launch commands. We can call those from a script in a layout. The demo scripts used below use `dbus`, take a look at the docs on scripting Konsole for details. I'm using the layout file `~/konsole_layouts/monitoring.json` for this example. This layout file represents two panes with one vertical split (horizontal orientation describes the panes being horizontally placed): `{ "Orientation": "Horizontal", "Widgets": [{ "SessionRestoreId": 0 }, { "SessionRestoreId": 0 }] }` Here's an example of a simple script using that layout, which will launch `fastfetch` in one pane and `btm` in the other:

```
#!/usr/bin/env bash # Define the commands to run cmd1="fastfetch" cmd2="btm" # Opens a konsole tab with a saved layout # Change the path to point to the layout file on your system # KPID is used for renaming tabs konsole --hold --layout $HOME/konsole_layouts/monitoring.json & KPID=$! # Short sleep to let the tab creation complete sleep 0.5 # Runs commands in Konsole panes service="$(qdbus | grep -B1 konsole | grep -v -- -- | sort -t"." -k2 -n | tail -n 1)" qdbus $service /Sessions/1 org.kde.konsole.Session.runCommand "${cmd1}" qdbus $service /Sessions/2 org.kde.konsole.Session.runCommand "${cmd2}" # Renames the tabs - optional qdbus org.kde.konsole-$KPID /Sessions/1 setTitle 1 'System Info' qdbus org.kde.konsole-$KPID /Sessions/2 setTitle 1 'System Monitor' What it does: Loads a layout with 2 panes, horizontally arranged Runs clear and then fastfetch in the left pane; runs btm in the right pane Wrap-up That's how you can accomplish opening a number of panes in konsole which run different commands. Using this kind of shortcut at the start of every work / programming session saved a little time every day which adds up over time. The marketing peeps would call it "maximizing efficiencies" or something. I hope some folks find this useful, and come up with many creative ways of making konsole work harder for them. Known issues and tips Running konsole from a Quick Command will open a new window, even if you want to just open a new tab. You may see this warning when using runCommand in your scripts. You can ignore it. I wasn't able to find documentation on what the concern is, exactly. The D-Bus methods sendText/runCommand were just used. There are security concerns about allowing these methods to be public. If desired, these methods can be changed to internal use only by re-compiling Konsole. This warning will only show once for this Konsole instance. Credits to inspirational sources Thanks to the Cool-Konsole-setup repo, where I found an example script for using commands in a layout via qdbus. Note: The scripts in that repo did not work as-is. This answer on Ask Ubuntu for improvements on the example scripts.
```

- [KStars v3.7.6 is Released](#) (2025/04/01 04:39)

KStars v3.7.6 is released on 2025.04.01 for Windows, MacOS & Linux. It's a bi-monthly bug-fix release with a couple of exciting features. Scheduler Plans VisualizedHy Murveit added a graph to the Scheduler page that displays visually the scheduler's plans--the same plans described in the log at the bottom of that page, and partially described in the scheduler's table. You can see altitude graphs for all the scheduler jobs, which are highlighted in green when that job is planned to be active. The next two nights of the plan can be accessed using buttons to the right (and left) of the graph. The graph can be enlarged or hidden by sliding the "splitter" handle above it up or down.PHD2 & Internal Guider

RMSMany users reported differences between the RMS value reported by Ekos internal guider vs PHD2. This is not a new issue as there was a difference in RMS calculations ever since Ekos Guider module was developed over a decade ago. In this release, we updated the internal Guider RMS calculations to use the same algorithm used by PHD2. This way, there is now a more consistent metric to judge the performance of the two guider systems.

Weather Scheduler IntegrationWeather station integration with the scheduler was improved. The weather enforcement is now global and not per job. If weather enforcement is enabled, you can adjust the Grace Period (default 10 minutes) in cases where the scheduler cannot be started due to a weather Alert or Warning. When a weather warning is received, existing jobs can continue to execute but new jobs will not be executed until the weather situation improves. Upon detecting a weather hazard, the scheduler execute a Soft shutdown mode where it can park the mount and dome, but still retains connection with INDI drivers to continue monitoring the weather situation. If the weather does not improve by the Grace Period, it then commences a full shutdown of the observatory. Otherwise, it should resume the job from where it was left.

Contrast Based FocusingJohn Evans added an option to allow focusing on non-star fields by using various contrast based algorithms. This is suitable for Lunar, Solar and planetary imaging.

Autofocus OptimizationJohn Evans added an option has been added to Focus that allows an Autofocus run to re-use a previous successful Autofocus result if the previous AF run occurred within a user-defined time period, say <10mins ago. This can speed up certain situations when using the Scheduler where multiple Autofocus requests can happen within a short period of time.

Imaging Planner ImprovementsHy Murveit pushed a new Imaging Planner catalog release along with improvements to the KStars Imaging Planner. It should now start up much more quickly on first use, or first use after a catalog upgrade. There were stability improvements. The catalog was extended to include 770 objects. Upgrade to KStars 3.7.6, use Data -> Download New Data to get the latest Imaging Planner catalog, and run Load Catalog in the Imaging Planner tool to take advantage of all these improvements.

Quick Go & RotateAdded support to Go and Rotate in Framing Assistant. This would command fast go to target and then followed by rotation to match position angle indicated. Simply adjust the Position Angle to your desired angle then command Ekos to solve and rotate in one go.

Scheduler Coordinates FlexibilityWolfgang Reissenberger introduced enhancements for handling target coordinates in the scheduler module: Add an option to switch the target coordinates between J2000 and JNow. This is interesting for those cases where the user wants to enter the coordinates manually, but has the coordinates only in JNow - for example when taking them over from the align module. Add a "use the current target" button. Currently, there is only an option to take over the current skymap center.

Furthermore, during the time where the moon is visible, it should be possible to schedule only those jobs that are not disturbed by moonlight (e.g. H-alpha captures). To enable this, a new optional constraint is introduced where the maximal moon altitude could be set.

Use PHD2-scheme graphToni Schriber modified the internal guider chart to use PHD2-scheme (RA/DEC) for graph of guide mount drift. This should help comparisons between PHD2 and internal guider more consistent.

- [SoK25 Highlights: My Season of KDE Journey With KEcoLab](#) (2025/04/01 00:00)

Hi everyone! I'm excited to share my experience so far as a mentee in the Season of KDE program. For those unfamiliar, Season of KDE is an amazing initiative by the KDE community that allows students and newcomers like me to contribute to open-source projects under the guidance of experienced mentors. This year, I'm working on the KDE Eco project, specifically creating comprehensive documentation—both written and video—for the KDE Eco Remote Eco Lab. This blog post is a chance for me to reflect on what I've accomplished, the challenges I've encountered, and my plans moving forward. What is the KEcoLab? The KDE Eco Remote Eco Lab is a project within the KDE Eco initiative, which is part of the KDE community's efforts to promote sustainability through energy-efficient Free Software. Specifically, the Remote Eco Lab provides a way for developers to measure the energy consumption of their software remotely, using a specialized lab located in Berlin. This lab was established with

support from KDAB. My role is to develop clear and accessible documentation to help users understand how to set up, use, and benefit from this tool. This includes written guides and video tutorials. I'm thrilled to be working on the video part of the project! What I've Done So Far Since starting the program, I've been diving into the project and making steady progress. Here's a rundown of what I've accomplished so far. First, there was research and familiarization. I began by exploring the Remote Eco Lab—reading any existing materials, studying its features, and understanding its purpose. I also had productive discussions with my mentor to align on goals and expectations. Second, work on written documentation. In the KDE Eco Remote Eco Lab project, we're working as a tight-knit team, with tasks divided among us to cover all bases. My main focus is on creating video documentation, but I also get to collaborate with my teammates who are tackling the written documentation. I often sit down with them to brainstorm, which has been a fantastic way to contribute beyond my primary role. It's exciting to see how our efforts—video and written—come together to make the project more accessible to users. Finally, video documentation prep. For the video documentation, I've been working on a script to guide users through the Remote Eco Lab's features. It's currently being refined with feedback from my mentors, who are helping me make it sharper and more user-friendly. I initially created a simple script to guide myself while making the video. As I progressed, I developed it into a much more detailed version to give others a clear understanding of the video's structure and flow. The mentors appreciated the detailed script and gave me their approval, as it provided them with a clear idea of how the video would take shape. You can refer to that detailed script here. To finalize how we're going to present the concept of reiterating until an energy drop is visible in our software, I've created a small proof-of-concept video. This video effectively demonstrates how, after reviewing your software reports, you can make the needed tweaks and then re-check the results using KEcoLab to decrease energy consumption and contribute to sustainability. Video: Proof of concept for the KEcoLab video documentation. (Video from Utkarsh Umre published under a CC-BY-SA-4.0 license.) While that's in progress, I've been digging deeper into the Remote Eco Lab itself—especially the energy consumption reports it generates, which are key for developers aiming to optimize their software. At the same time, I'm learning Kdenlive, an KDE's video editing tool, to bring the script to life. It's been a fun challenge to master, and I'm excited to create tutorials that will help users get the most out of KEcoLab. These steps have helped me build a solid foundation for the documentation, and I'm proud of the progress I've made! Challenges I've Faced Of course, the journey hasn't been without its hurdles. Here are a couple of challenges I've encountered: My work on the KDE Eco Remote Eco Lab hasn't been without hiccups. Initially, I struggled to understand the final energy consumption reports—the data felt overwhelming and confusing. With help from my mentors and some extra digging, I've started to get it, which is key for my video tutorials. Another challenge was setting up OBS Studio for recording—it kept crashing on my system. After some trial and error, including updating my drivers, I got it running smoothly. Beyond these, things have gone pretty well, and I'm learning a ton! While these challenges slowed me down at times, they've also been valuable learning opportunities. I'm growing more confident with each step! Outlook for the Future I'm honestly so excited to keep going with the KDE Eco Remote Eco Lab project. My big focus right now is the video documentation—I've been messing around with my script, getting some awesome feedback from my mentors, and I'm almost ready to hit record. I'm hoping to have a handful of tutorials done by the time Season of KDE wraps up. I'll be using OBS Studio to capture everything and Kdenlive to edit it into something that's easy to follow and actually looks good. My teammates are busy crushing it on the written guides, and I can't wait to see how it all comes together. Oh, and guess what? I just found out my talk got picked for the KDE India Conference 2025, happening April 4-6 in Gandhinagar! I'll be sharing my KDE Eco adventure and what I've been up to as a Season of KDE mentee. It's a little nerve-wracking but mostly thrilling, and I'm pumped to prep for it while juggling my video stuff! Final Thoughts Participating in Season of KDE has been an incredible experience so far. I'm not only sharpening my technical and creative skills but also getting

a front-row seat to the collaborative spirit of the KDE community. I'm grateful to my mentors Aakarsh, Kieryn, Karanjot, and Joseph for their guidance and to the community for this opportunity. I can't wait to see how the project evolves and to share my final update in the next blog post! Interested In Contributing? KEcoLab is hosted here. If you are interested in contributing, you can join the Matrix channels Measurement Lab Development and KDE Eco and introduce yourself. Thank you to the Season of KDE 2025 admin and mentorship team, the KDE e.V., and the incredible KDE community for supporting this project. Please feel free to contact me here: matrix id - @utkarshumre:matrix.org Email - utkarshumre@outlook.com linktree Thanks for reading!

- [Preparing KEcoLab Technical Documentation in Season of KDE '25](#) (2025/04/01 00:00)

Hey everyone!! Welcome to my blog post. I am Roopa Dharshini, a mentee in Season of KDE 2025 for the KEcoLab project. In this blog, I will explain my work in the SoK mentorship program. Getting Started With SoK For my proposal I crafted a detailed timeline for each week. With this detailed plan and with the help of my wonderful fellow contributors and mentors, I was able to complete all the work before the end of the mentorship program. CC-BY-SA-4.0 license.)" src="https://eco.kde.org/blog/images/2025-03-31-roopa-sok25-proposal.png" style="max-width: 100%; height: auto" /> I started by first week working to understanding the project's codebase, studying KEcoLab's handbook and existing documentation, setting up a GitLab wiki in the forked repository, and discussing the GitLab wiki's Merge Request (MR) feature. I explored and discussed various technical documentation tools with the mentors. Initially, we had planned to continue with GitLab, but later due to the flexibility of KDE's community wiki, we proceeded with that as our preferred documentation tool. CC-BY-SA-4.0 license.)"

src="https://eco.kde.org/blog/images/2025-03-31-roopa-sok25-usage-scenario.png" style="max-width: 100%; height: auto" /> I got to work creating an outline for the entire technical documentation. Usage scenarios scripts are essential for executing the automation pipeline in KEcolab. So, my fellow mentees and I started our documentation process with usage scenario scripting: we drafted a short page describing its importance, provided some scripts, and detailed their structure. This documentation is structured in a way that even non-technical contributors are able to follow the guidelines and create their own scripts. CC-BY-SA-4.0 license.)"

src="https://eco.kde.org/blog/images/2025-03-31-roopa-sok25-ci-cd.png" style="max-width: 100%; height: auto" /> After this, I wrote various texts for the technical documentation (CI/CD pipeline, Home Page) of the KEcoLab project. There was a change in the audience for our documentation: initially we focused on the users of KEcoLab, but later we decided to write documentation for both the people who wish to contribute and provide new changes to KEcoLab as well as those who use KEcoLab for their software measurements. This change had us writing in-depth technical documentation for developers who wish to change the code for better efficiency. The CI/CD pipeline is essential for the energy measurement automation in KEcoLab. Writing detailed CI/CD pipeline documentation that explains its use, structure, and job execution was challenging, yet rewarding. Final Documentation Links User Guide documentation for KEcoLab Users Usage Scenario Script documentation Accessing result documentation for users CI/CD pipeline documentation for contributors Contribution guidelines How did I apply to Season of KDE? CC-BY-SA-4.0 license.)" src="https://eco.kde.org/blog/images/2025-03-24-roopa-sok25-proposal.png" style="max-width: 100%; height: auto" />

Season of KDE is a mentorship program that happens every year between January and March. It is a three-month mentorship where mentees will be guided through a project they propose. You start by writing a proposal and timeline to work on from the projects listed on the KDE Ideas page. You tag the mentors in the issue, and they will review your proposal and check whether you are suitable or not. You can checkout my proposal for the KEcoLab project. After review, mentors will hopefully mark your proposal as accepted. And that's how I got into it! Challenges I faced Applying to SoK was not easy for me. I ran into my first challenge when I tried to create a new KDE Invent account. I thought

there were some technical issues with the website, so I tried every day to create an account (you are limited to one account creation chance per 24-hour period). After a long wait, I reached out to SoK admin Johnny for help, and he assisted me in creating an account. I was really scared to submit my proposal because there was only one week before the submission deadline, but I trusted my skills and submitted it. So, keep in mind that "it is never too late to apply." The second challenge was team collaboration. Similar to me, there were 2 other contributors selected for this project. I was brand new to KDE. At first it was hard to communicate with my other contributors, but later on we started to work really well together. Those are the main challenges I faced during my contributions to SoK. Challenges are never an end point; they are a stepping stone to move further. Thank You Note! Challenges make the journey worthwhile. Without any challenges, I wouldn't have known the perks of contributing to KDE in SoK. I take a moment here to thank my wonderful mentors Kiernyn, Aakarsh, Karanjot, and Joseph for guiding me throughout this journey. Also, I want to thank my fellow contributors to the project Shubhanshu and Utkarsh for collaborating with me to achieve what we proposed successfully. Finally, I am thankful to the KDE e.V. and the KDE community for supporting us new contributors to the amazing KDE project. KEcoLab is hosted on Invent. Are you interested in contributing? You can join the Matrix channels Measurement Lab Development and KDE Eco and introduce yourself. Thank you!

- [KDE Plasma 6.3.4, Bugfix Release for April](#) (2025/04/01 00:00)

Tuesday, 1 April 2025. Today KDE releases a bugfix update to KDE Plasma 6, versioned 6.3.4. Plasma 6.3 was released in February 2025 with many feature refinements and new modules to complete the desktop experience. This release adds three weeks' worth of new translations and fixes from KDE's contributors. The bugfixes are typically small but important and include: Fix glitch while scrolling with touch. Commit. ToolsAreaManager: Store windows as a vector. Commit. Fixes bug #501688 Kstyle: Don't replay scrollbar mouse event to same position. Commit. View full changelog

- [Wrapping Up My SoK Journey](#) (2025/03/31 18:49)

Introduction -Over the last 10 weeks, I had the opportunity to contribute to MankalaEngine by exploring and integrating new algorithms for gameplay, as well as working on adding the Pallanguli variant to the engine. My journey involved researching about various algorithms like Monte Carlo Tree Search (MCTS), implementing Q-learning, an ML-based approach, and evaluating their performance against the existing algorithms of MankalaEngine. Also assisted in reviewing the implementation of the Pallanguli variant. Implementing and Testing MCTS I first explored Monte Carlo Tree Search (MCTS) and implemented it in MankalaEngine. To assess its effectiveness, I tested it against the existing algorithms, such as Minimax and MTDf, which operate at depth 7 before each move. MCTS Performance Results - Player 1 Player 2 MCTS Win Rate Random MCTS 80% MCTS Random 60% Minimax MCTS 0% MCTS Minimax 0% MTDf MCTS 0% MCTS MTDf 0% The results was not good enough. This was expected because existing Minimax and MTDf algorithms are strong and operate at depth 7 before each move. Moving to Machine Learning: Implementing Q-Learning. Given MCTS's poor performance against strong agents, I explored Machine Learning (ML) techniques, specifically Q-Learning, a reinforcement learning algorithm. After learning its mechanics, I implemented and trained a Q-learning agent in MankalaEngine, testing it against existing algorithms. Q-Learning Performance Results - Player 1 Player 2 Q-Learning Win Rate Random Q-Learning 100% Q-Learning Random 98% Minimax Q-Learning 100% Q-Learning Minimax 0% MTDf Q-Learning 100% Q-Learning MTDf 10% Q-learning showed significant improvement, defeating existing algorithms in most cases. However, it still had weaknesses. Techniques Explored to Improve Q-Learning Results: To improve performance, I experimented with the following techniques: Using Epsilon decay to balance exploration (random moves) and exploitation (using learned strategies). Increased rewards for wins to reinforce successful strategies. Training Q-learning against

Minimax and MTDf rather than only against itself. Despite these improvements, Q-learning still could not consistently outperform all existing algorithms. After these experiments and research, I believe more advanced algorithms like DQN or Double DQN are needed to outperform all existing algorithms. This would also be an exciting project for this summer. Work related to Integration of Pallanguli Variant Apart from exploring ML algorithms, I also worked on integrating the Pallanguli variant of the Mancala game into MankalaEngine. My contributions included: Reviewing Srisharan's code, suggesting fixes and discussions. Creating Merge Request (MR) that allows users to input custom initial counters for Pallanguli. Conclusion -This journey has been an incredible learning experience, and I am grateful for the guidance of my mentors, Benson Muite and João Gouveia, who were always there to help. I look forward to continuing my contributions to the KDE Community, as I truly love the work being done here. Thank you to the KDE Community for this amazing opportunity!

- [On brightness, and calibrating your displays](#) (2025/03/31 17:50)

Many people are, understandably, confused about brightness levels in content creation and consumption - both for SDR and for HDR content. Even people that do content creation as their job sometimes get it really wrong. Why is there so much bad information about it out there? Before jumping into the actual topic, I want to emphasize that most people that have gaps in their knowledge about HDR and SDR are not to blame for it. The standards that define colorspace are usually confusingly written, many don't paint the full picture, finding the one you actually need can be difficult, some you need to pay for to even read, and generally there is not a lot of well organized and free information about this out there. When you have basically no information, you just go with what you do know - you see how Microsoft Windows does HDR for example, maybe you take a look at a draft for the sRGB specification or simply the Wikipedia pages, and do the best with what you have. The result is often less than ideal. Having worked on this stuff for a while now, and having read lots about it from people that actually know what they're doing, I think I know the topic well enough to clear up some misconceptions, but do keep in mind that my knowledge is limited too, and I may still make mistakes. If you're sure I got anything wrong, tell me about it! If you want an entry point for way more information than this blog post provides, check out [color-and-hdr](#). How brightness works with sRGB sRGB is the colorspace most content uses today. Despite that, very annoyingly, its specification is not openly available... but there's a draft version that you can download freely here, which is good enough for this topic. The (draft) specification defines two things that are important when it comes to brightness: a set of reference display conditions a set of reference viewing conditions (I'll call that "viewing environment" from here on) The reference display conditions are seemingly quite straight forward. The display luminance is 80cd/m², we have a whitepoint of D65, and a transfer function. Transfer functions describe how to calculate the output luminance from the encoded values of an image, and with sRGB that's $Y = X^{2.2}$ where Y is the relative luminance on the display, and X is the relative luminance on the input. The viewing environment has a few more parameters, but it's conceptually not difficult to understand: It describes how bright your environment is, what color temperature the lights in your room have, and how much your display reflects the environment at you. How to create sRGB content "correctly"? The assumption that many people take from the specification is that you should calibrate your display to 80cd/m². On its own, that information is completely wrong! It's obvious when you think about how end users actually view content: They set the brightness level of the display to what they're comfortable with in the current environment. You make the display really bright when you're outside, less bright when in a normally lit room, and even darker than that when the lights are off. The part that's missing with just calibrating the display to some luminance level is that you must take the viewing environment into account. Either you set up the sRGB reference viewing environment (with measurements!)... or you just don't. When you create content, in most cases you should do exactly the same thing as the person that will consume the content does: Just set the brightness to what's comfortable in the environment you're in. It still helps to keep your viewing

environment mostly fixed of course, lots of brightness changes mean you're constantly readjusting and that's not good. There's another big thing to take into account for sRGB, which is its confusing transfer function. The sRGB transfer function The sRGB specification doesn't just define a transfer function for the display, but it also defines a second transfer function. This sRGB piece-wise transfer function is if $X < 0.04045$: $Y = X / 12.92$ else: $Y = ((X + 0.055) / 1.055)^{2.4}$ and it's slightly different from gamma 2.2 in that it has that linear bit for the very dark area. The purpose of this transfer function is to optimize encoding of dark parts of the image - with 8 bits per color, gamma 2.2 becomes really small in the lowest few values. $1/255$ for example results in roughly 0.0000051 with gamma 2.2, and 0.0003035 with the sRGB piece-wise transfer function. This difference might sound insignificant, but it is noticeable. The most well known place of where the wrong transfer function is used is Microsoft Windows: When you enable HDR in Windows, it uses the piece-wise transfer function for sRGB content, instead of the gamma 2.2 transfer function that which your display uses in SDR mode. The result is that dark areas of SDR games and videos are brighter than they should be, and look "washed out". So when should you use the sRGB piece-wise transfer function? So far, I don't know of any case where you should, outside of working around that Windows problem in your application... I'm also only concerned with displaying images though, and not editing or creating them, so take that with a grain of salt. How brightness works with HDR Most HDR content uses the SMPTE ST 2084 transfer function. The specification for this is freely available here. SMPTE ST 2084 is a bit different from the sRGB spec, in that it only defines a transfer function but no complete colorspace or viewing environment. That transfer function is the Perceptual Quantizer (PQ): It tries to compress luminance levels in a way that matches how sensitive human eyes are in specific luminance ranges, and it's defined in absolute luminance - a PQ value of 0.0 means $\leq 0.005 \text{cd/m}^2$, and 1.0 maps to 10000cd/m^2 . The missing parts are defined by different specifications, rec.2100 and BT.2408. More specifically, rec.2100 uses the BT.2020 primaries with the PQ transfer function (or the HLG transfer function, but we'll ignore that here) and a recommended viewing environment for such HDR content: BT.2408 expands on that with an HDR reference white and graphics white, at 203cd/m^2 . This is mostly meant for the context of broadcasts, referring with "graphics" to logos or subtitles in the video stream. Despite the transfer function being "absolute", just like with sRGB, the luminance numbers don't mean anything in isolation. When displaying HDR content, just like with SDR, we need to take the viewing environment into account, and adjust luminance levels accordingly. How is this handled in Wayland? Every transfer function in the color management protocol has reference display conditions and a viewing environment attached to it, defined by a few parameters. Most relevant for this topic are a reference luminance, also known as HDR reference white, graphics white or SDR white minimum and maximum mastering luminances, basically how dark and bright the display the content was made for can go When content is displayed on the screen, the compositor translates between the viewing environment of the content, and the viewing environment of the user. While we don't usually have full knowledge of what exactly that viewing environment is like, the brightness slider in KDE Plasma provides a very good approximation by configuring the reference luminance to be used for content on the display. The calculation for this brightness adjustment is rather simple, in linear space you just do $\text{output} = \text{input} * \text{output_reference} / \text{input_reference}$ You can configure the maximum reference luminance (brightness slider at 100%) with the "Maximum SDR Brightness" in the display settings of Plasma 6.3. The minimum and maximum luminance your display can achieve can only be configured with the kscreen-doctor command line tool right now, but an easy to use calibration utility for this is nearly finished (and the default values are usually fine too). In general, this system is working really well... with one rather big exception. HDR in Windows games As mentioned before, Windows in HDR mode does sRGB wrong, but the story with HDR content is kind of worse. When you use Windows 11 on a desktop monitor and enable HDR, you get an "SDR content brightness" slider in the settings - treating HDR content as something completely separate that's somehow independent of the viewing environment, and that you cannot adjust the

brightness of. With laptop displays however, you get a normal brightness slider, which applies to both SDR and HDR content. The vast majority of Windows games expect the desktop monitor case: Static, never changing luminance levels, which are displayed on the screen without any adjustments whatsoever. Windows also didn't have a built-in HDR calibration tool until Windows 11, so nearly every Windows game ships with its own HDR calibration settings and completely ignores system settings. This doesn't just cause issues for Windows 11 laptops of course, but also for playing these same games with HDR on Linux. Until Plasma 6.2, we worked around that, also mostly not doing brightness adjustments, and the result was that those HDR calibration settings in games worked basically like on Windows. However, these workarounds broke Linux native applications that want to mix HDR and SDR in their own windows, made tone mapping worse, and blocked features like HDR on "SDR" laptop displays, so in Plasma 6.3 we had to drop them. This doesn't mean you can't play Windows games with HDR in 6.3 anymore, you just have to adjust their configuration to match the changed brightness levels. In most cases, this means you set the HDR paper white in games to 203cd/m², and then set the maximum luminance with the game's configuration screen, like this one from Baldur's Gate 3: [How to implement good HDR](#)

After ranting about how Windows games do it wrong, I should end this blog post by also explaining how to do it right. I will skip most of the implementation details, but on a high level if you're implementing HDR in a Wayland native application or toolkit, you should use the Wayland color management protocol get the capabilities of the compositor and/or graphics driver, specifically the transfer functions they support get the preferred image description from the compositor, and the luminances you're supposed to target from that. When using these luminance values, keep in mind that reference luminance adjustment the compositor will do! every time the preferred image description changes, get the new one and adjust your application to it now render for these parameters, and set the image description you actually ended up targeting on the surface, either through Vulkan or with the Wayland protocol (not both at the same time!) SDR things, like user interfaces in games, should use the reference luminance too if your application has some need to differentiate between "SDR" and "HDR" displays (to change the buffer format for example), you can do so by checking if the maximum mastering luminance is greater than the reference luminance now you can, and really should drop all HDR settings from your application. If HDR has a performance penalty in your application, a toggle to limit the app to SDR could still be useful, but everything else should be completely automatic and the user should not be bothered with calibration screens or similar annoyances

- [Akademy 2025 will be held in Berlin](#) (2025/03/31 09:57)

Akademy 2025 will be held at the Technische Universität Berlin (TU Berlin) in Berlin, Germany, from Saturday the 6th to Thursday the 11th of September. Akademy 2025 will be a hybrid event, combining on-site and remote sessions, and will include talks, workshops, Birds of a Feather (BoF) meetups, training and coding sessions. The conference is expected to draw hundreds of attendees from the global KDE community to discuss and plan the community's future and its technologies. Many participants from the broad free and open source software community, as well as local organizations and software companies, will also attend. The call for papers will open soon, and the registrations shortly after. We will soon update Akademy's website; in the meantime, follow us on Mastodon, Lemmy and Twitter to keep up to date with Akademy's news. About Berlin Berlin, the capital and largest city of Germany, has a population of about 3.7 million and is located in the country's northeastern part. With over 800 years of history, it has been a key center for politics, culture, and technology, serving as the capital of the Kingdom of Prussia and later the German Empire. The city is known for its cultural landmarks like the Brandenburg Gate and the Reichstag, along with renowned museums on Museum Island, a UNESCO World Heritage Site. Berlin is also a hub for innovation and art, hosting events such as the Berlin International Film Festival (Berlinale) and tech conferences like re:publica. It offers various recreational activities for residents and visitors. About the Technische

Universität Berlin (TU Berlin) Technische Universität Berlin (TU Berlin) is a public research university. is one of the 20 largest universities in Germany with around 35,000 students in over 100 degree programs. It is known for its strong engineering, computer science, and natural sciences programs. The university fosters innovation and international collaboration across disciplines. About Akademy For most of the year, KDE, one of the largest free and open software communities in the world, works online communicating over email, instant messaging, video-conferencing, forums and mailing lists. Akademy provides all KDE contributors with the opportunity to meet in person to foster social bonds, work on concrete technology issues, discuss new ideas, and reinforce the innovative, dynamic culture of KDE. Akademy brings together artists, designers, developers, translators, users, writers, sponsors and many other types of KDE contributors to celebrate the achievements of the past year and help determine the direction for the next year. Hands-on sessions offer the opportunity for intense work, bringing those plans to reality. The KDE community also welcomes companies building on KDE technology to Akademy, as well as those who are looking for opportunities.

- [February/March in KDE Itinerary](#) (2025/03/30 11:00)

In the past two months since the last update localization of KDE Itinerary has been improved, more ticket formats are supported and work on public transport information infrastructure continued, among many other things. New Features Localized units In locales using imperial units Itinerary now shows distance, speed and temperature values converted to the corresponding units. Distance and speed values in imperial units. If you are using one of those locales but still want to use metric units that of course remains possible, there's a switch on the settings page for forcing the use of metric units. Metric unit switch. Events We are in the middle of conference season, so there were several opportunities to watch Itinerary-adjacent talks: Felix, Jonah and Marcus presented Transitous at the Railway and Open Transport track of FOSDEM. Also at FOSDEM, Itinerary had a cameo-appearance in Hans-Jörg Happel's talk about Structured Email in the Modern Email track. I spoke at the Wikidata Data Reuse Days about Itinerary's use of Wikidata. Wikimedia Germany also covered that in a blog post. We had two presentations at the FOSSGIS Konferenz, with me introducing Transitous and Felix covering the routing engine behind it in more detail (both in German). Besides conference talks there were also a few related sprints and meetups: The bi-annual OSM Hack Weekend in Karlsruhe. Another iteration of the Carl's Kitchen Sprint. This continues with the Plasma (Mobile) Sprint in Graz three weeks from now. Infrastructure Work Transitous GBFS rollout Transitous has started to test first/last mile routing with shared vehicles in a few areas. If things hold up this is will be gradually added in more areas where there necessary data is available. Similar to the GTFS format for public transport schedule data, there's GBFS as the open standard for positions of currently available shared vehicles. That's covering anything from bikes over electric kick-scooters to cars. Client-side we are prepared for this already, KPublicTransport support this as routing parameters e.g. with OpenTripPlanner backends already. There's however still more work needed to expose this properly in the UI, not just as ad-hoc options as it's done now but also as part of a persisted personal routing profile. DOSIPAS ticket barcode container support The travel document extractor can now decode UIC's "DOSIPAS" ticket barcode container. That stands for "Double Signed Package Structure" and is the designated replacement for the UIC 918.3 container format. It uses the very compact but nasty to parse ASN.1 unaligned packed encoding rules (uPER) instead of zlib-compressed ASCII-ish content, and allows for more modern cryptographic signatures. Similar to UIC 918.3 it can contain multiple payloads, both standardized and vendor-specific ones. The predominant one we found so far is "FCB" (Flexible Content Barcode), an also ASN.1 uPER encoded common superset of all European ticket data models with several hundred (mostly optional) properties. We had two of the three FCB versions already implemented fortunately. DOSIPAS tickets are in use in some areas of France for regional trains, e.g. Grand Est and Normandie. While this generally should improve Itinerary's ability to detect tickets correctly, to the point of importing by just scanning the barcode even, there's a darker side to this as well, the "Double Signed" part in DOSIPAS.

That's (optional) infrastructure for shortlived ticket barcodes that the vendor app continuously regenerates, not unlike what 2FA apps do. The mechanism for this is documented, so we could also implement this of course. The challenging part here however is getting to the necessary secret key used to generate the dynamic signature. For now most of the DOSIPAS samples found in the wild are fortunately still static, e.g. in PDFs. Should the dynamic ones become mandatory at some point that would basically imply a mandatory use of the vendor app.

KPublicTransport journey subsection API Continuing the work around trip queries mentioned last time, KPublicTransport's journey API received to few changes to make selecting sub-sections of a journey easier and to make reassembling journeys from sub-sections possible. The main technical limitation for this so far was that the departure and arrival stops were treated specially and couldn't hold the same information as intermediate stops. That's a historical leftover from before we even had support for intermediate stops and meant that shortening/extending a journey would lose information. While maybe a seemingly small implementation detail this nevertheless required quite some effort, and will allow removing some limitations in e.g. how train trips can be edited in Itinerary. Fixes & Improvements Travel document extractor Added or improved travel document extractors for 12go, Amtrak, Color Lines, Eventyay, Flixbus, Ghotel, SBB, SNCF, UK national railways, Universe and VR. Fixed validity end date parsing in FCB customer card barcodes. Fixed a crash on VDV tickets without a basic ticket data block. All of this has been made possible thanks to your travel document donations! Public transport data Read operator and occupancy information from DB Zugportal onboard API. Fixed OJP journey queries using stop identifiers (used e.g. in Switzerland). Correctly merge journey sections with partial intermediate stop data. Itinerary app Correctly display coach/seat numbers when having a separate seat reservation or multiple travelers with different seat reservations. Don't show delay information for walking legs. Fixed showing intermediate stops for bus trips. Don't hide essential trip group map elements on low zoom levels. Fixed trip group map bounding box computations for trips without elements that change locations. Fixed data loss when realtime information would change the type of a reservation. Changed importing data from OSM to use Nominatim, as described here. Allow transfers to favorite location independent of the next reservation. Correctly handle nested events when determining transfers. Disabled Deutsche Bahn online ticket import as the corresponding API is no longer available. Don't show city name in location search results if that matches the station name. Update statistics automatically when removing trips. Fixed misrendered labels in wallet passes created by the vdv-pkpass converter. Itinerary also benefited from work on improving the Android platform integration of KDE apps as well as various fixes in the QtQuick Controls Breeze Style. 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.

- [Kaidan 0.12.2: Message Removal and Bubble Fixes](#) (2025/03/29 23:00)

Kaidan 0.12.2 fixes some bugs. Have a look at the changelog for more details. Changelog Bugfixes: Fix removing corrected message (melvo) Fix showing message bubble tail only for first message of sender (melvo) Download Source code (.tar.xz) (sig signed with 04EFAD0F7A4D9724) Linux (Flatpak on Flathub) Or install Kaidan for your distribution:

- [This Week in Plasma: zero VHI bugs and much more](#) (2025/03/29 04:00)

Welcome to a new issue of "This Week in Plasma"! Every week we cover the highlights of what's happening in the world of KDE Plasma and its associated apps like Discover, System Monitor, and more. This week we drilled into the outstanding bug lists, and drove the number of HI and VHI priority bugs down to their lowest ever numbers! In addition, we boosted performance, made high-visibility improvements to notification history and screen locking, implemented support for multiple cross-desktop standards, and way more! So, quite a big week. Notable UI Improvements Plasma 6.4.0 Notifications in the history popup now retain their interactive buttons, if they have any. If they don't but clicking on their

background normally does something when they're in pop-up form, then they'll show an "Activate" button when in the history. (Dominique Hummel and Kai Uwe Broulik, [link 1](#) and [link 2](#)) On the lock and login screens, The clock and interactive UI elements are now only shown on one screen at a time when using a multi-screen setup; they fade out on screens without the pointer or keyboard focus, leaving those screens free to display pretty wallpapers. (Yifan Zhu, [link 1](#) and [link 2](#)) System Settings' Display Configuration page gained some UI Improvements; now the screen arrangement view is hidden when there's only one screen, and with more than one, there's a big obvious screen chooser at the top of the page to make it clear which screen is selected, and when there are any disabled but connected screens. (Oliver Beard, [link](#)) Improved the appearance of the Comics widget when it hasn't been set up with any comics yet, or when there's been an error of some kind. (Christoph Wolk, [link 1](#) and [link 2](#)) You can now also use Meta+Tab and Meta+Shift+Tab to switch between windows, in addition to the current shortcuts. This supports our push to have all global actions include the Meta key for at least one of their shortcuts. (Vlad Zahorodnii, [link](#)) KWin's "Fade Desktop" virtual desktop switching effect now has a customizable duration. (Konstantin Kharlamov, [link](#)) Plasma's Notifications now respect requests to play sounds using "sound hints". (Ruslan Khabibullin, [link](#)) The Breeze cursor theme with dark cursors is now named "Breeze Dark", and vice versa for the ones with light cursors. (Niccolò Venerandi, [link](#)) Notable Bug Fixes Plasma 6.3.4 Fixed the most common Plasma crash! This one could happen when unplugging screens, especially with a dock involved in the process somewhere. This was the final VHI priority bug! (David Edmundson, [link](#)) Fixed the remaining causes of two common KWin crashes. (Xaver Hugl, [link 1](#) and [link 2](#)) Fixed a severe crash in Breeze-themed apps that we accidentally introduced in Plasma 6.3.3 alongside a change to fix a bug with color scheme support for creative color schemes. The change itself was fine, but it exposed a pre-existing issue that was also benign on its own. When the two combined... kaboom. But no more, now that it's fixed! (Albert Astals Cid, [link](#)) Fixed a clipboard bug that caused non-ASCII text from items re-ordered in the history to become mangled when pasted. (Fushan Wen, [link](#)) Fixed a bug that made it impossible to remove previously-added languages on System Settings' Region and Language page. (Christoph Wolk, [link](#)) Fixed a bug that made Discover sometimes fail to show the current version of an app or package being updated to a newer one. (Ismael Asensio, [link](#)) Fixed a visual glitch affecting auto-hiding top-positioned fit-content panels. (Niccolò Venerandi, [link](#)) Plasma 6.4.0 Fixed several issues in the desktop & wallpaper settings window that prevented certain pages from being scrollable when the content was long. (Christoph Wolk, [link 1](#) and [link 2](#)) Notifications created by apps using the Notifications portal that specify something to happen when clicked now actually perform that action. (Kylie CT, [link](#)) Fixed the System Tray's adherence to the part of the StatusNotifierItem spec that allows tray icons to ask to display a context menu on left-click. (Kai Uwe Broulik, [link](#)) Flatpak apps from non-standard user-defined repositories are now shown on System Settings' Flatpak Permissions page. (Harald Sitter, [link](#)) Fixed a bug that broke scrolling on certain scrollable views in the Application Dashboard widget. (Tomislav Pap, [link](#)) When using the "Choose Player Automatically" feature of the Media Player widget, the actual name of the player is now shown on the lock screen (when using the feature to show media information on the lock screen) instead of the text "Choose Player Automatically". (Fushan Wen, [link](#)) Fixed a bug that prevented the Media Frame widget from pausing the slideshow when hovered with the pointer, as it was originally intended to do. (Christoph Wolk, [link](#)) Other bug information of note: Zero very high priority Plasma bugs! (3 last week). Current list of bugs 18 15-minute Plasma bugs (21 last week). Current list of bugs Notable in Performance & Technical Plasma 6.4.0 Massively improved performance when making screen recordings in Spectacle using the VP9 video format, which is used by default. (Arjen Hiemstra, [link](#)) Implemented support for the wp_fifo_v1 Wayland protocol. (Xaver Hugl, [link](#)) Implemented support for the Clipboard Portal. (David Redondo, [link](#)) Made KWin more resilient against the issue of windows moving to strange positions when changing the screen arrangement or number of screens. (Xaver Hugl, [link](#)) Plasma's log output is pretty quiet now thanks

to Christoph's hard work. Nonetheless, he continued that work to help get us to zero! (Christoph Wolk, [link 1](#), [link 2](#), [link 3](#), [link 4](#), [link 5](#), [link 6](#), [link 7](#), [link 8](#), [link 9](#), and [link 10](#)) 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. You can help KDE by becoming an active community member and getting involved somehow. 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: Triage and confirm bug reports, maybe even identify their root cause Contribute designs for wallpapers, icons, and app interfaces Design and maintain websites Translate user interface text items into your own language Promote KDE in your local community ...And a ton more things! You can also help us by making a donation! Any monetary contribution — however small — will help us 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, feel free to push a commit to the relevant merge request on invent.kde.org.

- [The Open Source Promotion Program](#) (2025/03/29 00:00)

The KDE community will again participate in the Open Source Promotion Plan (OSPP), a program in which students can contribute to open source projects. Burgess Chang, is the KDE community contact. As part of OSPP 2024, Hànyáng Zhāng (张汉阳) added Android support to Blinken. The work done is described in a series of blog posts, available in both English and Mandarin. The Android version is available from the KDE F-Droid nightly repository. Unlike the Google Summer of Code, where stipends are funded by a company, stipends are primarily funded by the Chinese government with options for open source communities to contribute additional stipends if they wish to have more students participate in their projects than they get allocated. It is good that there is recognition that contributing to open source software is a skill that students should acquire. The range of contributions that can be made in OSPP is not just limited to programming, contributions to other aspects that improve the open source software ecosystem such as translation and documentation are welcome. As it is a government funded program, there is a little more oversight to ensure tax payer funds are well spent. In particular, for most projects, contributions should be made to a publicly available repository associated with the project and that student participants are selected primarily based on their project application. The plan aims to increase the programming and software engineering skills of students by encouraging them to participate in real world projects during their vacation period. While it is funded by the Chinese people, open source projects with contributors from all over the world apply to participate, and students from any part of the world can also apply to participate. Mandarin and English are the official communication languages for the program, knowledge of one of these is sufficient to participate in the program. The OSPP website lists the dates for each phase of the program. Important dates for this year are: 04 April - 04 May: Project submission period for approved open source communities 09 May - 09 June: Student project application period 01 July - 30 September: Coding and development period for accepted projects

- [Kaidan 0.12.1: Voice Message and Password Change Fixes](#) (2025/03/28 23:00)

Kaidan 0.12.1 fixes some bugs. Have a look at the changelog for more details. Changelog Bugfixes: Do not highlight unpinned chats when pinned chat is moved (melvo) Fix deleting/sending voice messages (melvo) Fix crash during login (melvo) Fix opening chat again after going back to chat list on narrow window (melvo) Increase tool bar height to fix avatar not being recognizable (melvo) Fix width of search bar above chat list to take available space while showing all buttons (melvo) Fix storing changed password (melvo) Fix setting custom host/port for account registration (melvo) Fix crash on chat removal (fazevedo) Move device switching options into account details to fix long credentials not being shown and login QR code being temporarily visible on opening dialog (melvo) Allow setting new password on error to fix not being able to log in after changing password via other device (melvo) Download Source code (.tar.xz) (sig signed with 04EFAD0F7A4D9724) Linux (Flatpak on Flathub) Or install

Kaidan for your distribution:

- [Wrapping Up Season of KDE'25](#) (2025/03/28 19:34)
- [Web Review, Week 2025-13](#) (2025/03/28 13:05)

Let's go for my web review for the week 2025-13. OpenAI's Studio Ghibli meme factory is an insult to art itself Tags: tech, ai, machine-learning, gpt, politics, culture, art, copyright Sure, a filter which turns pictures into something with the Ghibli style looks cute. But make no mistake, it has utter political motives. They need a distraction from their problems and it's yet another way to breach a boundary. Unfortunately I expect people will comply and use the feature with enthusiasm... <https://www.bloodinthemachine.com/p/openai-studio-ghibli-meme-factory> Trapping misbehaving bots in an AI Labyrinth Tags: tech, ai, machine-learning, gpt, security When a big player has to prepare a labyrinth of AI generated content to trap bots used to feed generative AI learning pipelines... something feels wrong. <https://blog.cloudflare.com/ai-labyrinth/> Improved ways to operate a rude crawler Tags: tech, ai, machine-learning, gpt, copilot, satire Don't underestimate how much of a skill making a stupid crawler can be... https://www.marginalia.nu/log/a_115_rude_crawler/ Proof of work reverse proxy to protect against scrapers Tags: tech, ai, machine-learning, gpt, copilot, security And yet another reverse proxy to use as a scraper deterrent... It looks like several are popping every week lately. <https://git.sr.ht/~runxiyu/powxy> Exploring Generative AI - The role of developer skills in agentic coding Tags: tech, ai, machine-learning, copilot, ide, tools, programming Again that confirms that all the hype and grand announcements are not deserved. It also gives a good idea of the skills which are required to use those tools, clearly the setup process is involved if you want to don't want to be overwhelmed and drowning in bad code. <https://martinfowler.com/articles/exploring-gen-ai.html#memo-13> Scallop, a Language for Neurosymbolic Programming Tags: tech, ai, machine-learning, neural-networks, logic, prolog This is definitely an interesting declarative language. Looking forward to more such neurosymbolic approaches. <https://www.scallop-lang.org/> Servo vs Ladybird Tags: tech, web, browser, foss A good look at both incumbents in the web browser engine space. Still quite some way to go but the results are interesting already. <https://thelibre.news/servo-vs-ladybird/> REST in Peace? Django's Framework Problem Tags: tech, python, django, rest, sustainability, community There's a sustainability issue for the REST support with Django. Hopefully this will resolve. <https://danlamanna.com/posts/rest-in-peace-djangos-framework-problem/> git-who: Git blame for file trees Tags: tech, version-control, git, tools Looks like a neat little tool to explore git repositories. <https://github.com/sinclairtarget/git-who> Quadlet: Running Podman containers under systemd Tags: tech, tools, containers, podman, systemd Looks like a nice way to orchestrate rootless podman containers. <https://mo8it.com/blog/quadlet/> Closing the chapter on OpenH264 - Pixels Tags: tech, video, codec, patents, foss Or why software patents can get in the way... You can work around them somehow, but that quickly leads to shipping binaries you can't properly check. <https://bbhtt.space/posts/closing-the-chapter-on-openh264/> Things that go wrong with disk IO Tags: tech, io, storage, filesystem, databases A reminder that writing on disks is a longer process than you could suspect. Many things can go wrong on that chain. <https://notes.eatonphil.com/2025-03-27-things-that-go-wrong-with-disk-io.html> C++/Rust Interoperability Problem Statement Tags: tech, rust, c++, interoperability Looks like there's movement at the Rust Foundation level to have better C++ and Rust interoperability. We'll see what comes to fruition, this could be interesting. It's needed for sure. <https://github.com/rustfoundation/interop-initiative> Use the rr debugger without HW performance counters ! Tags: tech, debugging, tools Interesting fork of rr to have time travel debugging with software counters. Hopefully will allow using rr in environments where it's limited by lack of access to hardware performance counters. <https://github.com/sidkshatriya/rr.soft> Postel's Law and the Three Ring Circus Tags: tech, foss, protocols, design, standard Nice post about the practical impacts of Postel's law. It's especially problematic in the case of Open Source software. Companies producing proprietary software even use that to their advantage.

<https://alexgaynor.net/2025/mar/25/postels-law-and-the-three-ring-circus/> War story: the hardest bug I ever debugged Tags: tech, debugging, web, browser, google Interesting story... when you end up turning to v8 having a bug in the field, you're really in trouble.

<https://www.clientserver.dev/p/war-story-the-hardest-bug-i-ever> Why developers question everything - Tim Hårek Tags: tech, craftsmanship, programming, estimates, risk, complexity Or why analogies with physical work don't work...

<https://timharek.no/blog/why-developers-question-everything/> Sun Tzu wouldn't like the cybersecurity industry Tags: tech, security It's better if you prepare your security policies properly... <https://kellyshortridge.com/blog/posts/sun-tzu-wouldnt-like-the-cybersecurity-industry/> The Worst Programmer I Know Tags: tech, team, productivity Trying to measure individual productivity is definitely a trap. You'd better not try, otherwise you'll have wrong behaviors or you'll punish the wrong persons. <https://dannorth.net/the-worst-programmer/> How to Write Blog Posts that Developers Read Tags: tech, blog, writing A bit cynical at times, but shows tricks to improve the writing and style of blog posts. If I ever find the time to write something sizeable again I guess I'll try some of them. <https://refactoringenglish.com/chapters/write-blog-posts-developers-read/>

Teach to Learn: Why Sharing What You Know Makes You Smarter Tags: teaching, learning I like this attitude obviously... Go out and teach! Share what you learn! <https://hardmodefirst.xyz/teach-to-learn-why-sharing-what-you-know-makes-you-smarter> Post Apocalyptic Computing Tags: tech, low-tech, history, reliability Interesting rambling and exploration. What would a computer built to last a century look like? <https://thomashunter.name/posts/2025-03-23-post-apocalyptic-computing> Bye for now!

- [KDE Snap updates, Kubuntu Beta testing, Life updates!](#) (2025/03/27 19:20)

Help us Beta test Kubuntu Plucky Puffin! Kubuntu Plucky Puffin (25.04) Beta released Kubuntu work: Fixed an issue in apparmor preventing QT6 webengine applications from starting. Beta testing! KDE Snaps: Updated Qt6 to 6.8.2 Updated Kf6 6.11.0 Rolling out 25.04 RC applications! You can find them in the -candidate channel! Life: I have decided to strike out on my own. I can't take any more rejections! Honestly, I don't blame them, I wouldn't want a one armed engineer either. However, I have persevered and accomplished quite a bit with my one arm! So I have decided to take a leap of faith and with your support for open source work and a resurrected side gig of web development I will survive. If you can help sponsor my work, anything at all, even a dollar! I would be eternally grateful. I have several methods to do so: GoFundMe Patreon Github Donorbox If you want your cool application packaged in a variety of formats please contact me! If you want focused help with an annoying bug, please contact me! Contact me for any and all kinds of help, if I can't do it, I will say so. Do you need web work? Someone to maintain your website? I can do that too! Portfolio Thank you all for your support in this new adventure!

- [Kubuntu Plucky Puffin \(25.04\) Beta released](#) (2025/03/27 18:30)

The beta of Kubuntu Plucky Puffin (to become 25.04 in April) has now been released, and is available for download. This milestone features images for Kubuntu and other Ubuntu flavours. Pre-releases of Kubuntu Plucky Puffin are not recommended for: Anyone needing a stable system Regular users who are not aware of pre-release issues Anyone in a production environment with data or workflows that need to be reliable They are, however, recommended for: Regular users who want to help us test by finding, reporting, and/or fixing bugs Kubuntu, KDE, and Qt developers Other Ubuntu flavour developers The Beta includes some software updates that are ready for broader testing. However, it is an early set of images, so you should expect some bugs. Highlights include an update to KDE Plasma 6.3. We STRONGLY advise testers to read the Kubuntu 25.04 Beta release notes before installing, and in particular the section on 'Known issues'. You can also find more information about the entire 25.04 release (base, kernel, graphics etc) in the main Ubuntu Beta release notes and announcement.

- [Twinimation Studios Presents Their Next Krita Workshop: Becoming an Artist on a Budget](#) (2025/03/27 00:00)

Twinimation Studios have released a new Krita workshop, and we wanted to give them a chance to introduce their new offering to Krita's users: Greetings everyone! Entering the art world is sometimes seen as an expensive endeavor. From art schools to subscription based software, artists across different fields tend to have notable expenses. But have you ever wondered if you can become an artist without spending a fortune? Twinimtion Studios is back to answer the question with our very first full workshop! Becoming an Artist on a Budget is a specialty made guide guide to help aspiring artists begin their artistic journey WITHOUT breaking the bank. This workshop consists of 9 main videos bundled into one easy to digest package, along with some special bonus showcase videos as well. Included is also a bonus freebie list of numerous artistic products ideas to begin a paid art hobby or career. There should have been a video here but your browser does not seem to support it. Within this workshop, we provide tips and tricks on how one can begin their art journey for completely free. After reviewing a list of affordable resources to learn art skills, we recommend numerous free art programs with a special spotlight on Krita! We explain how versatile Krita is, and how it can be used across numerous different art fields, such as animation, comics, and painting! Following some other drawing tutorials, the workshop concludes with a special lesson on entrepreneurship, where we explain how aspiring artists can create a paid hobby or full business through their artwork while remaining on a budget. With so many people wanting to enter the art scene and build a career from it, we hope this workshop will be a helpful guide for all of those who wish to create their own artistic brand. Additionally, we have many other Krita focused animation courses on our website! Twinimation Studios was founded by instructors Andria and Arneisha Jackson; MFA graduates who've studied animation for 9 years and want to share their professional knowledge with the world. We provide tutorials on different styles of animation, character design, illustration, film creation and so much more! Look forward to our future tutorials and workshops where we will continue to expand our repertoire to fit several different art fields. Here is a link to the workshop: Become an Artist on a Budget

- [Week 3-4 />](#) (2025/03/26 18:30)

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**

