

Trakr.

Projects contain tickets. That's the whole system.

A self-hosted issue tracker and work-management platform built to be unceremonial — and to include what the incumbent makes you buy separately. A query language, a customer support desk, ITSM and ITIL service management, SLA policies, timesheets and billing, ten reports where every figure clicks through, on-premise AI, and a pipeline that takes a ticket to a pull request.

Track work. Ship it. Skip the ceremony.

Self-hosted

14 languages

On-premise AI

10

REPORTS, ALL DRILLABLE

25+

AI PROVIDERS

148

IN-APP HELP ARTICLES

14

INTERFACE LANGUAGES

THE PROBLEM

The tracker became the work.

Somewhere between the first board and the third year, the tool that was supposed to record work started generating it. Fifteen ticket types. A configuration nobody dares touch. Page loads measured in seconds. A support desk, a service desk, a timesheet tool and a reporting add-on each on their own licence — and the whole estate hosted somewhere you did not choose, at a per-user price that rises every time you hire.

Configuration debt

A workflow so elaborate that changing it requires a specialist and a change window.

Add-on sprawl

The support desk, the service desk, the timesheets and the reports are four more invoices.

Per-seat gravity

The cost of letting one more person see the board is a reason not to.

AI you cannot use

The assistant is in someone else's cloud, so it is off in every regulated environment.

THE SOLUTION

One application, one database, and a data model you can hold in your head.

Simple by default

Projects contain tickets. Tickets have statuses. Methodology is a *lens* that adds optional views — Kanban, Scrum, Waterfall or SAFe — and never adds fields. The default is none of them.

Complete when needed

Support desk, service management, SLA policies, timesheets, billing, budgets, automation, reporting and git are in the product, not in a marketplace.

Yours to run

Self-hosted on bare metal, in containers, or on Kubernetes. Every asset, font, library and help article is bundled — the server works with no internet at all.

THREE DECISIONS WORTH READING

A merge that deletes nothing

Merging duplicates re-points comments, work logs and attachments so the authors and timestamps stay original. The absorbed ticket keeps its key, its history and its place in the customer portal — because deleting it would strand the customer who filed it.

Old keys resolve forever

Move a ticket between projects and it is re-keyed — but the old key keeps working, permanently, and keys are never re-used. Bookmarks, pasted links and email threads survive, so a customer's reply cannot fork the conversation into a duplicate.

Reordering is not editing

Dragging the backlog writes exactly one row and does not touch the ticket's updated date — so a morning spent sorting does not destroy the stale-ticket report, the daily digest and every “changed in the last day” query in the organisation.

TRACK & FIND

The tracker, the query language, the service desk.

1 Core tracking

- Types, priorities, statuses and services are **configuration rows, not fixed lists** — each organisation, and optionally each project, defines its own with colour and icon
- Eleven link types plus parent and subtask relationships
- Workflows with transitions and enforced conditions, rendered as a state diagram — and a refused board drag links straight to the diagram that refused it
- Move between projects with full re-keying and remapping; merge duplicates with a preview that takes the same request body as the apply
- **Attachment versioning** — replace a file and the previous version stays downloadable, with restore; inline preview uses a deliberately narrower allow-list than download
- Ticket templates at user, project and organisation scope, optionally bound to a type so the right dialog opens pre-filled
- Bulk update, transition and link — with values resolved by name, so one call edits a selection spanning projects
- List, board, roadmap, calendar, timeline, backlog and sprint views; cycles as sprints, phases or program increments
- **Delegations** — “I’m away until the 19th” redirects assignment and makes routing skip you

3 Support desk & service management

- Three project worlds — standard, helpdesk and service — separated **server-side**, not by whichever page is asking
- A public customer portal with emailed guest access codes, internal versus public comments, canned responses and per-project settings
- The support desk **refuses to be switched off** while projects still use it, naming them
- ITSM and ITIL as opt-in switches. **Everything is a ticket** — service request, incident, problem and change are kinds, not separate records, so comments, work logs, search, bulk edit, merge, SLA, automation and the API all apply unchanged
- Problems carry root cause, workaround and a known-error flag; changes carry type, risk, planned window, implementation and backout plans and an outcome
- Approvals enforced by the workflow condition, not by an endpoint — a review board is simply an approval asked of a group
- SLA policies with per-priority targets, business calendars and holidays, pause-on-status and breach notification; a running timer keeps the deadlines it was started under
- Email-to-ticket across five providers with reply threading by ticket key

2 The query language

- A purpose-built query language over core fields, people, dates, collections, time, links, full text and SLA state
- Fourteen operators; ordering comparisons on priority use your organisation’s own rank, so “priority at least high” means what it says
- **Historical operators** — was, was not, changed, changed from, changed to — replayed from the activity log
- Current-user, period and relative-date functions
- Design guarantees that matter: a negation reaches the empty rows, a bare date means the whole day, every sort is settled so paging is deterministic, and **a clause is never quietly ignored** — an unparseable value is refused with the clause quoted back
- An AI mode that returns **a query, never results** — so paging, columns, grouping, export and bulk selection keep working, and the query stays visible, editable and savable
- Autocomplete, recent searches, saved filters, export and a per-user column chooser

4 Time, timesheets & money

- Work logs in familiar duration notation with a billable flag — **no approval workflow, no sign-off, no locking period**. Honest logging only
- A personal timesheet as a day-tile grid *including the empty days* — the blank Wednesday is the only actionable thing on the page
- Exports as CSV, spreadsheet and PDF, including a multi-month single PDF for an arbitrary set of months
- Statistics spanning up to ten years with day through year buckets and year-over-year comparison, reporting both person-days and full-time-equivalent days and labelling which is which
- Administrators may view and file time on another person’s timesheet — filed under the subject, with the activity feed recording who did it
- Daily limits for regular and overtime minutes, resolved per field down organisation, tenant and platform, enforced on every write path
- Billing and budget reporting in mandays at a price per manday

REPORT, AUTOMATE, BUILD

Where a tracker becomes a management tool.

5 Reporting

- Ten reports: ad-hoc generator, billing, budget, SLA, resolution metrics, time in status, team performance, statistics, timesheets and AI spend
- A personal dashboard with eighteen widget types, and shareable dashboards
- **Time in status stores nothing** — the whole report replays the activity log, so it can never drift; it reports the median as well as the mean, because a mean is what one forgotten ticket ruins
- Team performance in three lenses — workload as of today, delivery, and effort — and it **deliberately refuses a composite productivity score**, because a blended number rewards closing small tickets and cannot be argued with by the person it describes
- Utilisation reads *null*, with a link to the setting, when no working day has been defined — rather than a percentage of an invented eight hours
- **Every figure drills through** to the tickets behind it, as an ordinary search — so paging, columns, export and bulk selection all work on the drilled list, and the query is left in the bar to be read and saved
- Scheduled delivery daily, weekly or monthly on rolling windows only — with organisation policy on external recipients re-checked *at delivery*, not merely on save

7 AI, agents & code

- **Over twenty-five providers, sixteen of them self-hostable** — the whole AI feature set runs on-premise with nothing leaving the network
- Twelve AI features across search, suggestion, translation, generation, merging, templates, assistance, bulk editing and reporting
- Cost control: tiered routing with budget-aware downshift, a price book, rate limits, a circuit breaker and a spend report
- A **Model Context Protocol server** with nine tools, eight resources, five prompts and revocable per-organisation keys, so agents operate the tracker natively
- **Ticket to pull request**: AI validation produces refined acceptance criteria and blocking issues, then one of seven command-line coding agents runs on a branch under a per-run and per-day budget cap and opens the request
- An integrated git server with repository browser, releases, pull requests and an in-browser editor — plus a two-way issue mirror with six external providers where the sync loop is broken by stored state, not memory
- Import from CSV and from the incumbent's CSV and XML exports

6 Automation

- Rules at platform, tenant, organisation and project scope, as a union rather than a ladder
- Triggered by a schedule, by events, or both — a frequency and a time, never a cron expression somebody can mistype into silence
- Twelve quality checks plus an open-ended condition builder over twenty-six fields and twenty-four operators, constrained by field kind so an impossible pairing is refused at save
- Nineteen actions from setting fields and transitioning through assignment, labelling, linking, commenting, creating tickets and calling webhooks
- A once mode that turns a rule into **recurring work**, so recurring tickets need no feature of their own
- Guarantees: a rule-authored write never triggers another rule; the workflow is never overruled; one failing action costs only that action; dry run performs nothing and reports what it would do
- Assignment routing with round-robin or least-busy modes over member, group, skill or named pools, pre-chosen fallbacks, and a live preview listing every candidate and why each is or is not eligible

8 Administration

- Guided organisation and project setup wizards that leave nothing behind if abandoned
- **Organisation restructuring** — a drag-and-drop board for moving, merging and splitting, rewriting ownership across roughly thirty tables in one transaction, with collisions requiring an explicit resolution rather than a guess
- **Configuration transfer** — export an organisation or a project as one portable bundle and import it elsewhere. Nothing in a bundle is an identifier, no people and no secrets ever travel, and an import never deletes
- Eighteen email templates resolved per field down four scopes, with a deliberately tiny placeholder language and a refusal to publish a customer subject that drops the ticket key
- Announcements, audit forwarding to ten sinks per organisation over signed HTTP or syslog, backups, self-update, migration management and nightly housekeeping
- A cluster with front-end, back-end and scheduler roles and **exactly one designated backup node** that takes over within ninety seconds, without a restart and without waking anyone

WHO BUYS IT

Anyone paying four invoices for one workflow.

Software & product teams

The primary fit. Tickets, boards, git bindings, pull requests, releases, coding agents and a query language — with an import path off the incumbent.

Digital agencies

Budgets in mandays at a price per manday, a billing report, timesheets, and per-client projects under one organisation. Built for billing by the day.

IT service desks

Full ITSM and ITIL: services, incidents, problems, known errors and changes with risk, windows and backout plans — with approvals enforced by the workflow.

Outsourced support

Helpdesk projects, a branded customer portal with emailed guest access, canned responses, and inbound mail with reliable reply threading.

Managed service providers

Tenant and organisation multi-tenancy, per-tenant payment accounts, white-label branding down to the custom domain, and plan limits you define.

Regulated organisations

Public sector, defence suppliers, healthcare and finance: fully self-hostable, air-gap friendly, encrypted secrets, audit forwarding — and AI that runs entirely inside the network.

Enterprises with an identity estate

Azure AD, Google Workspace, LDAP and Active Directory, SAML, SCIM provisioning and time-based multi-factor — with per-organisation configuration loaded at runtime.

Teams leaving the incumbent

Two migration paths: import the export, or run the **issue mirror** against your existing tracker — provably reversible, because disconnecting leaves both sides intact.

European mid-market

Fourteen interface languages including Luxembourgish, euro-denominated budgets, and data that never leaves the continent unless you send it there.

WHO USES IT

Platform superadmin

Organisation admin

Manager

Member

Contributor

Viewer

Customer

Seven ranked roles, reused for project membership. Two are worth noting: the **manager** gets the analytical reports without needing administrator rights, because a manager who cannot see who is overloaded is being asked to manage blindfolded — and the **customer** is excluded by design from every people picker, assignee list, headcount and timesheet, with no path anywhere that lets one log work.

WHAT IT CONSOLIDATES

USUALLY A SEPARATE LICENCE	IN TRAKR
Issue tracker	Core
Customer support desk	Included
IT service management	Included
SLA management	Included
Timesheets & billing	Included
Reporting add-on	Included
Automation add-on	Included
Git hosting	Included

TECHNICAL SPECIFICATION

Fast and functional beats slow and pretty.

Runtime	Java 25 on Spring Boot 4; native image supported for sub-second startup with no virtual machine in the runtime image
Database	PostgreSQL, with the ticket table partitioned by project, full-text indexing and materialised views
Front end	Server-rendered with vanilla ES6 — no framework, no package manager, no build step ; rich text, board drag-and-drop and date pickers all hand-written
Assets	Everything vendored — zero content-delivery-network references anywhere
API	REST with an OpenAPI document and an explorer offering samples in eleven languages, including COBOL, Fortran and RPG
Deployment	Bare metal with a sandboxed service unit, containers or Kubernetes; splits into front-end, back-end and scheduler nodes with a hot standby
Scale	1 423 source files, 117 entities, 929 test files including 35 browser tests

SECURITY & GOVERNANCE

- Five parallel authentication mechanisms — session-validated, genuinely revocable tokens, form login, per-organisation single sign-on loaded at runtime, provisioning tokens and agent keys — plus LDAP, Active Directory and SAML
- **Time-based multi-factor** with encrypted secrets and full audit
- Twelve-character minimum passwords with complexity, a blocklist and a breach check; sliding session renewal re-reads the account, so a downgrade bites at once
- Tenant isolation through six independent mechanisms, with a resource in another organisation answering **not found, never forbidden** — confirming existence is itself an answer
- **AES-256-GCM** for every stored secret; no secret is ever read back
- Per-request content-security nonces with no inline scripting, strict transport security with preload, and rate limiting per address *and* per account
- Outbound request validation re-checked **at delivery** — a hostname that resolved publicly an hour ago can resolve to loopback now
- Audit across eight categories and thirty-eight actions, streamed on export and forwarded
- User deletion asks where the work goes before writing anything, and offers anonymisation — credentials wiped, sessions revoked, content never destroyed

COMMERCIAL

Licensed per deployment. Never per seat.

Plans are data, not a decision the product makes for you: you define the tiers, the limits and the prices, and one switch disables billing entirely for an on-premise install. Licensed per deployment or per tenant, with unlimited users.

REVENUE	PER YEAR	REVENUE	PER YEAR
Up to €500k	€900	€100M	€13,000
€1M	€1,500	€500M	€37,000
€5M	€2,500	€1B	€52,000
€25M	€5,500		

Annual licence on your own turnover, ex VAT. The schedule is marginal like income tax — each slice of turnover is charged at its own rate, so there are no cliffs — with a €900 minimum and any yearly increase capped at 25%. Every tranche carries unlimited users, projects and organisations inside one tenant, self-hosting and support that rises with the tranche. Hosting and hardware are yours, AI runs on your own provider keys and is billed by them, and consultancy, setup, migration and training are quoted separately. Illustrative: a tracker at €8 per user a month bills a two-hundred-person organisation €19,200 a year; at €5M turnover this is €2,500.

NEXT STEP

Mirror one project.

Leave your current tracker running. Point the mirror at one repository and work in both for a fortnight. Disconnecting changes nothing on either side.

Sales	tim@automatize.eu
Partners	luc@automatize.eu
Web	automatize.eu

Trakr is not affiliated with or endorsed by the vendors of any tracker it can import from or mirror. Specifications may change.