

PRODUCT AND PROJECT DELIVERY IN ONE WORKSPACE

PlanMagnet

The roadmap, the backlog and the sprint that ships it are the same records, in one place.

ROADMAPS

BOARDS AND SPRINTS

FEEDBACK INBOX

RELEASES

DELIVERY OPERATIONS



01 The short version

PlanMagnet is one workspace where the plan and the work are the same records, so the roadmap you show leadership is built out of the sprint your team is actually running.

Most teams do not have a planning problem so much as a plumbing problem. The roadmap lives in one tool, the tickets in another, the specs in a third, and the feedback from the people who asked for the work in a spreadsheet nobody opens. Every week somebody spends an afternoon making those four things agree with each other, and by Friday they disagree again.

PlanMagnet removes the plumbing by removing the separate tools. An initiative on the roadmap points at the epic that delivers it. The epic breaks into the tasks on the board. The board runs the sprint. The sprint feeds the velocity. The release bundles what shipped, and the request that started it all is still attached at the other end of the chain. Portfolios, programmes, risks, quality gates, budgets and timesheets sit in the same workspace rather than in add-ons bought separately.

It is honest about where it is. PlanMagnet is pre-launch. The connected delivery workspace, the feedback inbox, versioned releases, the delivery-operations suite and one-click import from the tools you are on today are all built. A spec editor, an objectives framework, burndown charting and most of the AI assistance are specified and committed but not yet shipped, and this document says so every time it matters.

01 The promise, in numbers

THE PROMISE

One workspace, one data model. The plan and the work stop being two systems that have to be reconciled.



1

One connected model

Roadmaps, backlogs, sprints, releases and delivery operations are one data model with one identity and one audit trail, not five subscriptions kept in step by hand.

5

Ways in, and back out

Import and export both ship for Jira, Linear, Asana, GitHub and Trello, so adopting PlanMagnet is a reversible decision rather than a one-way door.

21

Governed resource types

From portfolios down to timesheets, every part of the product is a named resource with its own permission level, seeded into every workspace automatically.

Free

Viewers, on every tier

Only people who create or change work are billed. Executives, clients and commenters read, follow and vote at no cost on every plan.

Who it is for

Product teams

Requests arrive from everywhere and die in a spreadsheet. A feedback inbox with voting and computed priority scoring sits next to the roadmap it feeds and the release that answers it, so prioritisation is evidence you can show rather than an opinion you have to defend.

Digital agencies

Multiple clients, each with their own methodology, grouped under portfolios. Budgets and expenses per project, time logged against real tasks, and a free read-only seat you can safely hand a client.

Engineering teams

Configurable Kanban and Scrum boards, sprints with start, complete and cancel, story points, velocity and explicit blocking relationships. The tracker is not a separate purchase from the plan it is delivering against.

Consulting firms

Engagements modelled as projects with custom fields and workflows per engagement type, billable time rolling into timesheets that go through a recorded submit and approve cycle rather than an email thread.

Delivery leaders and executives

One roadmap across teams, portfolios and programmes for the shape of a growing org, and live dashboards drawn from real items instead of a status deck somebody assembles the night before.

Distributed and remote-first teams

Discussion, decisions and dependencies travel with the work item instead of scrolling out of a chat channel, so somebody three time zones away can pick up the context without a meeting.

What breaks today

None of these are exotic failures. They are the ordinary, weekly cost of running delivery across tools that were never designed to know about each other.

The roadmap and the work live in different tools

The roadmap is accurate on the day it is presented and drifting by the following Tuesday, because nothing connects an initiative to the tasks that move it.

- A roadmap item points directly at the epic, project and milestone delivering it, so the plan references real work rather than describing it from memory.

Feedback arrives everywhere and counts nowhere

Requests land in support tickets, sales calls and hallway conversations, and the loudest single voice wins because nobody merged the four ways people described the same thing.

- A feedback inbox collects requests as real records with vote counts, a computed priority score and links to the initiative that would address them.

Teams outgrow simple trackers but dread enterprise weight

You either stay on a tool that cannot express a sprint, or buy one that takes a quarter to configure and an administrator to keep running.

- Boards, columns, swimlanes, work-in-progress limits, sprints and custom workflows configure per project, so the process fits the team rather than the other way round.

Delivery operations are bought as add-ons

Risk registers, quality gates, budgets and time tracking end up in three more subscriptions or, more often, in spreadsheets that nobody trusts by month three.

- Portfolios, programmes, a risk register, quality gates, budgets and expenses, time tracking and timesheets all ship as first-class parts of the same workspace.

Decisions vanish into chat

Six months later somebody asks why the design changed and the answer is in a thread nobody can find, so the team relitigates a settled question.

- Threaded comments, mentions and decisions attach to the item they are about, and every change is attributable and recorded.

Nobody can defend a date

The answer to when does it land is a guess with confident body language, which is remembered as a commitment by everyone in the room.

- Velocity is computed across completed sprints from real items, and the trailing average is presented as exactly what it is: arithmetic you can check, not a prediction.

Switching tools is a project in itself

Teams stay on a tool they have outgrown because migration means an export, a consultant and a fortnight of lost history.

- Import and export both ship for Jira, Linear, Asana, GitHub and Trello, so a trial can start with your real backlog and can be reversed if it does not work out.

Access is all-or-nothing

Contractors and clients see budgets and hourly rates because the tool has one notion of member, and the alternative is a read-only export nobody maintains.

- Every action is checked against a role held per resource type, so a client stakeholder can follow the roadmap and vote on requests while budgets, expenses and timesheets stay invisible to them.

04 What you get

PlanMagnet is assembled from a small number of connected building blocks. Each is a real record type with its own permissions, history and links to the others.

PLAN

Roadmaps, initiatives and milestones

Group initiatives by quarter, theme or objective, and link each one down to the epic, project and milestone that delivers it. Timeline and Now, Next, Later views read off the same records. Automatic status roll-up from task movement is specified but not yet built.

DELIVER

Boards, sprints and backlogs

Kanban and Scrum boards with configurable columns, swimlanes and work-in-progress limits. Sprints start, complete or cancel, carry story points, and produce velocity and point-based capacity. Burndown charting is on the roadmap.

TRACK

Work items, links and dependencies

Epics, stories, tasks, bugs and subtasks with custom types, custom fields and workflow states. Parent and child links and explicit blocking relationships are real records, so a chain of blockers can be walked rather than guessed at.

04 What you get

CONTINUED

LISTEN

Feedback inbox and competitive context

Feature requests carry vote counts and a computed priority score, and can be linked both to the competitor whose move prompted them and to the roadmap initiative that would answer them.

SHIP

Versioned releases

Bundle epics and fixes into a version with markdown release notes, track readiness, and tie the release back to the items that went into it. Auto-generated release notes are planned.

GOVERN

Risk register and quality gates

Risks scored on a likelihood by impact matrix, each with an owner, a status and a mitigation. Quality gates are named checkpoints a release must pass, reviewed as an explicit operation that records a verdict.

SCALE

Portfolios and programmes

Products at the top, the projects delivering them grouped beneath, and cross-project dependencies made visible, so multi-team planning keeps a shape as the org grows.

ACCOUNT

Budgets, expenses, time and timesheets

Budgets with an approval step, expenses with category and receipt, time logged against the task it was spent on, and timesheets moving through submit, approve or reject with the sign-off recorded.

AUTOMATE

Rules, templates and custom fields

No-code trigger and action rules with an execution log, project templates that materialise a configured project in one action, and workspace-defined custom fields whose values attach to individual records.

05 Put to work

These are illustrative scenarios describing what the product is built to do. PlanMagnet is pre-launch, so none of them is a customer story, and none carries a metric that was not measured.

HEAD OF PRODUCT

A roadmap the whole org can read, that stays true

Initiatives are grouped the way the organisation thinks, by quarter, theme or objective, and each one links down to the epic and tasks that deliver it. Timeline and Now, Next, Later views talk to executives and engineers in the language each prefers, off the same records.

- One link replaces the weekly reconciliation, and the roadmap references real work instead of describing it.

SCRUM MASTER

Run the sprint on a board that fits the team

Pull from a prioritised backlog into a sprint, then work it on a board with the columns, swimlanes and work-in-progress limits the team chose. Moving work is a drag, not a form. At sprint close, incomplete work carries over cleanly and velocity is computed from what actually finished.

- The board reflects reality because updating it is cheaper than not updating it.

PRODUCT OWNER

Keep a backlog that is ready, not a graveyard

Triage incoming work, estimate and rank it, and split oversized epics into stories sized to fit a sprint. Bulk editing and saved filters make grooming a focused session. Priority is explicit and visible, so the team pulls from the top rather than debating what is next.

- Sprint planning starts from a list that is genuinely ready to build.

ENGINEER

Every issue keeps the shape it needs

Custom issue types carry their own fields and workflow states, so a bug moves through triage, in progress, in review and done while a task follows a path that fits it. Assignment, labels, links and dependencies keep related work connected, and every change is recorded.

- Reopening a regression six months later starts from a history rather than from archaeology.

PRODUCT MANAGER

Turn scattered feedback into a ranked signal

Requests become records in one inbox with vote counts and a computed priority score. Each can be linked to the initiative that would address it, and to the competitor whose move makes it urgent, so prioritisation weighs demand and market together.

- A roadmap you can defend with evidence, and a way to close the loop with the people who asked.

RELEASE MANAGER

Plan a release and tell everyone it shipped

Group the relevant epics and fixes into a versioned release, track readiness, and require named quality gates to pass before it goes out. Release notes are written in markdown against the version, and the release stays tied to the items that went into it.

- Launching becomes a coordinated event with an auditable answer to is this ready.

PORTFOLIO LEAD

Give multi-team work a shape that survives growth

Organise work as products at the top level with the projects that deliver each one grouped beneath, inside portfolios and programmes. Cross-project dependencies become visible records, so a delay shows its ripple everywhere it matters.

- Leadership reads one picture of every product and the work serving it, with no separate portfolio spreadsheet.

DELIVERY LEAD

Catch risks and gate quality before you ship

Risks are logged deliberately with a likelihood, an impact, an owner and a mitigation, and bucket onto a matrix where the top-right corner is what deserves attention now. Quality gates are named checkpoints with criteria, reviewed as an explicit operation that records approved or sent back.

- Is this ready to ship has an attributable answer on the record instead of a nod in a corridor.

PROGRAMME MANAGER

Log time and approve it without spreadsheets

Time is logged against the task it was spent on, so a work item accrues its real time-spent. At period end those entries roll into a timesheet the person submits, and an approver accepts or rejects it. Approving is a separate, permission-gated action, and who signed off and when is recorded on the timesheet.

- Effort sits beside estimate on the same record, and approved periods stay trustworthy after the fact.

FINANCE-FACING LEAD

Keep every project on budget while it is still in flight

A project carries a budget with an amount and a currency, and expenses are logged against it with a category, a date and a receipt. Setting a budget and having it approved are separate, audited actions, so a plan of record exists before spend accrues. The over or under figure is shown, not quietly clamped.

- Burn is readable per project while a course correction is still cheap.

OPERATIONS LEAD

Spin up new projects from proven templates

Take a project that worked, with its boards, columns, statuses and starter work, and save it as a template. Applying that template materialises a complete configured project in one action, alongside the methodology starters that ship with the product. Templates are editable records a team refines as its process improves.

- Every engagement starts from your best practice rather than from nothing.

OPERATIONS LEAD

Automate the rules the team keeps forgetting

Compose rules from triggers and actions with no code: route incoming work, keep a parent epic in step with its children, escalate a stale item, or post an update. Rules are visible and editable records with an execution log, so the team can see exactly why something moved. Runs are metered by plan.

- Process becomes infrastructure rather than a discipline that decays.

ENGINEERING MANAGER

Run delivery on data rather than gut feel

Dashboards read velocity, cycle time, throughput and cumulative flow off the team's real items rather than a parallel spreadsheet. A widening cycle time flags a bottleneck while it is still cheap to fix, and a steady velocity makes a forecast something you can show your working for.

- The report is always current because the data is the work, with no reconciliation step in between.

DELIVERY LEAD

Move off your current tracker without losing the history

Import brings projects and issues in from Jira, Linear, Asana, GitHub or Trello, and export sends them back out again. A new project is created against the methodology the team actually runs, and adopts the matching workflows and statuses from the first task.

- Adoption becomes a reversible experiment you can start with your real backlog.

WORKSPACE ADMIN

One workspace, many teams, still governed

Six roles are seeded into every workspace automatically, from a delivery owner who holds everything down to a stakeholder viewer safe to hand an external client. Each role holds a level per resource type, and a resource simply omitted from a role is a resource that role cannot reach.

- A contractor sees their project, a client sees a read-only roadmap, and nobody sees a budget by accident.

PRODUCT MANAGER

Bring competitive context into prioritisation

Competitors are records that link to the feature requests and roadmap items they bear on. When a request comes up for prioritisation, the competitive angle sits next to the vote count and the priority score, and unlinking is as easy as linking when the landscape shifts.


-
- A strategically important gap stops losing to whichever request happened to collect the most votes.

06 Built AI-native

AI that plans from your own sprints, and admits what it has not built


Most planning tools bolt a model onto a tracker that can see your tickets but not the initiative they deliver, nor the velocity that would make a date credible. Grounding is the whole game in planning, and grounding is exactly what a bolted-on assistant does not have. PlanMagnet is AI-native in the only sense that matters here: the roadmap, the backlog, the sprint history, the dependency graph and the delivery metrics are one set of records, so a suggestion can point at the work behind it rather than at a summary of a summary.

That structural choice buys a narrow, deliberate scope. PlanMagnet commits to four planning decisions — draft work from an idea, summarise a thread into the decision, forecast a date from real throughput, and flag at-risk work — and builds them one at a time rather than shipping a chat box and calling it intelligence. One of the four is live in the service today: a description goes in, and a validated list of typed, prioritised work-item suggestions comes back. The other three are committed roadmap, and this catalogue marks each one as such.



The governance is the part that is finished. Every AI call routes through a single primitive that cannot complete without writing an append-only record first: who asked, what triggered it, which provider and model answered, tokens in and out, latency, the parsed decision and the outcome. Failures are recorded, including the case where the provider answered but the answer could not be parsed. Only a short non-identifying summary of the request is stored, never the raw text. And the live capability has no ability to write to your backlog at all, so nothing lands unreviewed as a matter of structure rather than policy.

06 Four structural commitments



One place to reason from

Roadmap, backlog, sprints, blockers and delivery metrics are the same records. Nothing has to be reconciled against an outside tracker before a planning question can be answered.




It suggests, it never writes

The live capability hands back a list and nothing else. What you keep is created through the ordinary flow with its usual checks, numbering and audit trail. There is no privileged AI route into your plan.



Every call leaves a record

One append-only line per call, written before the answer is returned: model, timing, decision, success or failure. A call that errors still leaves a line rather than disappearing.



It only sees what you can see

AI requests are permission-checked exactly like every other action. An external agent driving PlanMagnet inherits the rights of the account you connected it with, and never more.

What the AI can actually do

Work-item breakdown

Paste a written feature request and receive a short list of typed, prioritised work items with optional descriptions and estimates.

Live today. It returns suggestions only and has no ability to create anything itself.

Choice of model per request

Gemini, OpenAI and Anthropic are all wired in, each with a documented default. Name one for a single request and it is honoured exactly.

Live today. Leave it unset and PlanMagnet works down its list, falling back on failure; calls are time-bounded.

Validated output

Every suggestion is measured against the item types and priority levels PlanMagnet actually uses, capped in number, and rejected outright if it has no title.

Live today. A garbled answer produces a clear error, never a half-list or an invented item.

Unavailable rather than invented

A provider whose key is missing reports itself unavailable at call time instead of degrading into a plausible-sounding answer.

Live today. Failing loudly is the only safe behaviour for a planning tool.

Append-only call record

One row per AI call capturing workspace, acting identity, trigger, provider, model, token counts, latency, the parsed decision and status.

Live today. Written before the result returns, so the attempt exists even when the call does not.

Dependency graph traversal

Blocking and blocked-by are real links between items rather than a convention in a comment, so a chain can be walked to where a delay starts.

Live today. It is the substrate an at-risk call needs to name the blocker instead of emitting a score.

Trailing velocity

Points per completed sprint, the trailing average, cumulative flow and cycle time, computed from your own finished sprints.

Live today, and labelled honestly as arithmetic you can check rather than a prediction.

Typed planning tools for agents

A published Model Context Protocol server exposes projects, work items, sprints, boards, milestones, dependencies, budgets, risks, releases, quality gates, roadmaps, time tracking and automations as typed tools.

Live today. It is currently PlanMagnet's most complete AI surface, and it sits outside the app.

Reading the call record back

Listing and opening the recorded AI calls, and an in-app history view over them.

In progress. The rows accumulate correctly today, but nothing yet exposes them for reading.

Confirmation before an agent deletes

A confirmation step and a dry-run mode on the destructive tools an external agent can reach.

In progress. Permissions still bound what an agent can touch, but within those rights it can delete faster than a person would.

Grounded delivery forecast

A delivery window for a project, sprint, roadmap item or release, with a confidence signal and references naming the exact sprints and velocity values behind it.

Planned. With no completed sprint history it is designed to say insufficient history rather than invent a date.

Planned, and marked as planned

At-risk work with a stated reason

Flags for stale, blocked and over-committed items, each carrying a reason in plain words rather than an opaque score.

Planned. It is designed to route a flag into the risk register only on human confirmation.

Semantic feedback de-duplication

Connect requests that share no keywords, propose duplicate clusters for a human to accept, and merge them with votes consolidated.

Planned. The correctness requirement is that one person asking twice must still count once.

Thread summarisation

Condense a long discussion into what was decided and what happens next.

Planned, and designed to summarise inside the workspace so message bodies do not travel out just to produce a summary.

And what keeps it in bounds

No unreviewed writes

The live capability has no code path that creates a work item. What you confirm is created through the ordinary flow, with its usual permission checks, key assignment and audit trail. This is a structural property, not a setting that can be relaxed.

Permission-checked at the edge

AI requests are authorised the same way every other action is, before they reach any resolver. An agent operating PlanMagnet inherits the rights of the account it was connected with, and cannot exceed them.

Failures recorded, not swallowed

A provider call that throws still writes a record with the error and the latency. A provider that answers unparseably is recorded as a provider success with the parse failure captured and no decision attached, and the caller gets an explicit error.

Prompt summaries, never raw text

The record stores a short non-identifying summary of what was asked rather than the request itself, so the audit trail does not quietly become a second copy of your content.

06 And what keeps it in bounds

CONTINUED

Keys resolved at call time

Provider credentials are never held in the product's source. Each adapter resolves its key at the moment of the call from secure storage, and a missing key fails as provider unavailable.

Output bounded by your own vocabulary

Suggestions are accepted only if their type and priority are ones PlanMagnet genuinely uses and their estimate is a finite, non-negative number. The list is capped so a single call cannot flood a review queue.

Explicit model selection

Naming a model pins that request to it, and an unrecognised name fails plainly rather than silently substituting a different one. Automatic fallback only happens when you did not choose.

06 And what keeps it in bounds

CONTINUED

Workspace scoping is non-negotiable

Planned retrieval features are specified so that one workspace's corpus can never influence another workspace's suggestions, and every AI record is indexed by workspace and time.

Spend pre-flight on background work

The planned scheduled sweeps are designed with a per-workspace off switch and a spend check before they run, so a background loop cannot quietly consume a workspace's AI budget.

Metering and caps, stated as missing

Token counts and latency are recorded today; a per-call cost figure and a per-workspace AI credit allowance are not. Both are planned, and until they land there is no in-product spend cap.

06 Reachable from, and in numbers

REACHABLE FROM

- The published Model Context Protocol server — the most complete AI surface today, driving PlanMagnet from the agent you already work in
- The federated GraphQL API, where the live breakdown capability is reachable as one permission-checked call
- The command-line tool, for scripted and pipeline use
- The Node and Python SDKs, over the same API surface
- The editor extension, bringing assigned work into the sidebar — planned
- An in-app assistant that knows which screen you are on — planned, and honestly absent today
- An AI operations screen for reviewing suggestions and auditing what ran — planned
- Scheduled planning sweeps on a daily, mid-sprint and weekly cadence — planned

IN NUMBERS

- Four planning decisions are committed in the specification; one of them is live in the service today
- Three model providers are wired in — Gemini, OpenAI and Anthropic — each with a documented default model
- One append-only record is written per AI call, before the result is returned to the caller
- Provider errors and unparseable responses both produce records; neither is discarded
- The live capability has zero write paths into the backlog; creation happens only through the ordinary flow
- 155 typed planning tools across 18 areas of the product are exposed to external agents, alongside the shared platform tools
- 24 of those tools are destructive and currently have no confirmation step or dry-run mode, which is stated rather than hidden
- Token counts and latency are recorded on every call; per-call cost is not yet recorded

06 What is not built yet

Three of the four committed AI decisions are not built. There is no delivery forecast with a confidence signal, no at-risk sweep, and no thread summarisation in the product today, and the semantic de-duplication of feedback is specified rather than shipped. There is no button for the live capability inside the app — it is reachable through the API, the command-line tool and agent tooling, and the app's own AI screen is a placeholder. The call records accumulate correctly but cannot yet be read back, so an administrator cannot browse them. Per-call cost is not recorded and there is no AI credit cap. Bringing your own model key is not offered. And around two dozen destructive tools reachable by an external agent still have no confirmation step, so until that lands the real control is how narrowly you scope the account you hand an agent.

07 Who does what

Six roles are seeded into every workspace automatically, and each holds a level per resource type rather than a single global rank. A resource omitted from a role is a resource that role cannot reach at all — the omission is the access control.

Delivery Owner

Owns the whole footprint

Accountable for how the workspace runs PlanMagnet end to end: methodology and field configuration, integrations, automation, project-finance approvals, and the audit surface. This is the only role that withholds nothing, and the only one that can act on the moderation queue rather than merely read the workspace quota and automation audit.

- Configure methodology and custom fields
- Approve budgets and timesheets
- Own integrations and automation rules
- Act on the moderation queue
- Read the full audit surface
- Hold every resource at full control

Portfolio Lead

Runs delivery across teams

The programme-office or agency-operations lead behind the budget and timesheet approval flows, running portfolios, programmes and projects across teams and configuring methodology, templates and integrations for all of them. Deliberately edits rather than owns the day-to-day backlog and feedback inbox, and reads the audit surface rather than acting on it.

- Structure portfolios and programmes
- Approve budgets and timesheets
- Own templates and custom fields
- Configure external syncs
- Own boards, sprints and milestones
- Read workspace quota and automation audit

Product Manager

Owns the product spine

Owns the product spine — the roadmap and its initiatives, the feedback inbox with its votes and competitor context, and the releases that ship them — working the backlog that delivers them. Has no reach into budgets, expenses, timesheets, automation or integrations, because finance and payroll are not part of the job.

- Shape roadmaps and initiatives
- Triage the feedback inbox
- Ship versioned releases
- Link competitors to requests
- Work epics, stories and dependencies
- No access to any financial record

Quality Lead

Governs readiness

Keeps the risk register honest and signs off the quality gates a release has to pass, with the ability to edit the work and links a gate decision depends on. Reads the wider portfolio and feedback context but cannot reshape the roadmap it is signing off against, and sees no financial data.

- Own the risk register
- Score likelihood and impact
- Record a quality-gate verdict
- Hold or release a version
- Edit the work a gate depends on
- Read-only on the roadmap

Delivery Contributor

Does the work

The everyday engineer, designer or consultant seat. Moves work across boards and sprints, maintains links, dependencies and custom field values, raises and votes on feature requests, and logs billable time into a timesheet they submit — but cannot approve their own timesheet or see a budget.

- Work items across boards and sprints
- Maintain links and dependencies
- Set custom field values
- Raise and vote on requests
- Log time and submit a timesheet
- Cannot approve or open a budget

Stakeholder Viewer

Watches, safely

The read-only seat designed to be handed to an executive or an agency client. Follows the roadmap, portfolio and delivery status and votes on feature requests, while every commercially sensitive surface — budgets, expenses, timesheets, integration configuration and the audit trail — is omitted entirely.

- Follow the roadmap and portfolio
- Read delivery status
- Vote on feature requests
- No budgets or expenses
- No timesheets or rates
- No integration configuration

08 Under the hood

PlanMagnet is a workspace-native product built on a shared multi-tenant platform. It contributes the product and project management domain — projects, work items, boards, sprints, roadmaps, releases, portfolios, risks, budgets and the rest — and inherits identity, tenancy, permissions, billing, files, conversations, notifications and audit from platform services that already exist rather than reimplementing them per product.

The API is a federated GraphQL schema served through public gateways. Every query and mutation field carries an explicit authorisation directive naming the resource, the action and the scope it requires. The gateway reads those directives from the composed schema and allows or denies before a request ever reaches a resolver, and the decision fails closed: if nothing grants the action, it is denied. Product services never validate tokens and never re-implement permission checks, so authorisation cannot drift between one module and another.

Transactional data sits in a relational store with graph relationships available for traversing dependency and link structures, with row-level policies isolating each workspace. Cache and event streaming are shared platform infrastructure, and the product publishes and consumes domain events so cross-product workflows can react to delivery changes without direct service-to-service calls.



THE SPLIT THAT MATTERS

Product logic is owned by PlanMagnet; identity, tenancy, authorisation, billing, files and audit are owned by the platform underneath it. That is why a governed workspace with role-based access, single sign-on, directory sync and an audit trail is available from the free tier rather than gated behind an enterprise plan — it is not a PlanMagnet feature, it is the ground the product stands on.

08 Capability by area

Domain model

Portfolios, programmes, projects, roadmaps and roadmap items, milestones, releases, work items with parent and child links and explicit dependencies, boards and columns, sprints, feature requests, competitors, risks, quality gates, templates, custom fields, budgets and expenses, time entries and timesheets, automation rules and external syncs.

Tenancy isolation

Every workspace, project and item is scoped to its tenant, with row-level policies enforcing the boundary in the data store rather than relying on query hygiene alone.

Authorisation

21 product-owned resource types with three ordered levels. Each action falls into a read, write or privileged tier that decides which levels may perform it, so delete, approve and review are structurally separated from everyday editing.

Agile analytics

Velocity across completed sprints, point-based sprint capacity, cumulative flow and cycle time computed from real items. Burndown charting is on the roadmap.

08 Capability by area

CONTINUED

Migration

Import and export both ship for Jira, Linear, Asana, GitHub and Trello, plus issue import from exported files. Connector depth continues to mature ahead of launch.

AI primitive

A single call path every AI capability must route through, which cannot complete without writing an append-only record first, with multi-provider selection and strict output validation on top of it.

Automation

No-code trigger and action rules with an execution log, metered against plan quotas. The action execution engine is still being completed.

Events

Domain events published on delivery changes and consumed from platform services, so cross-product workflows react without direct service-to-service coupling.

08 Standards it speaks, tools it plugs into

OPEN STANDARDS IT SPEAKS

- GraphQL, as a federated schema with a published contract
- Model Context Protocol, for typed tool access by external agents
- SAML-based single sign-on
- SCIM directory provisioning
- TLS for all client and API traffic
- Markdown for release notes and long-form descriptions
- CSV and PDF export for reporting
- Scoped API keys for programmatic access
- Multi-currency amounts on budgets, expenses and billable rates
- Webhooks for outbound events — planned

PLUGS INTO

- Jira — import and export
- Linear — import and export
- Asana — import and export
- GitHub — import and export, plus commit and pull-request linking
- Trello — import and export
- Source control and continuous integration, so a deploy can advance a release
- Chat, so updates land where the team already talks
- Calendar, for ceremonies, milestones and sprint dates
- Gemini, OpenAI and Anthropic as interchangeable model providers
- Workforce capacity planning through a sibling product — planned

08 How it deploys

Multi-tenant cloud

The default. A workspace on shared, regionally hosted infrastructure with tenant isolation enforced in the data layer, available from the free tier.

Regional hosting

Data is hosted in compliance-grade cloud regions, with data-residency options discussed as part of an enterprise agreement.

Dedicated tenant or self-host

Offered on the enterprise tier for organisations whose compliance or residency requirements rule out shared infrastructure.

Web-first, by choice

The full product runs in the browser today. Desktop and native mobile companions are planned rather than shipped, and this catalogue does not pretend otherwise.

Programmatic access

The same API that serves the app serves the command-line tool, the SDKs and the agent tool surface, so automation is never a second-class path.

09 Security & trust

PlanMagnet does not reinvent security per product. It inherits a platform security model and adds a domain-specific role catalogue on top, which is the reason a small team gets governance a small team would not otherwise build.

Encryption in transit and at rest

All traffic between a client and the API is encrypted with TLS, and data is encrypted at rest in compliance-grade infrastructure. Secrets are encrypted and access-controlled rather than held in application source.

Fail-closed by default

If nothing in the caller's roles grants the action, it is denied. Access is granted by explicit declaration; there is no permissive default to forget to turn off.

Authorisation enforced at one point

Permission checks happen at the API gateway, never inside the product service. Every query and mutation carries a directive naming the resource, action and scope it requires, and the gateway resolves it from the composed schema before a resolver runs.

Identity, never raw credentials

Requests reach the product carrying a verified actor identity rather than a token, so the product cannot mis-validate a credential because it never sees one.

Multi-tenant isolation

Every workspace, project and item is scoped to its tenant, with row-level policies enforcing the boundary at the data layer as well as the query layer.

Level per resource, not per user

A role grants an ordered level on each of 21 resource types. Actions are tiered so that delete, approve and review are structurally unavailable to editing-level holders.

Role catalogue seeded automatically

Six workspace roles plus organisation and billing presets are seeded into every workspace, organisation and billing account, so a new workspace is governed on day one rather than after somebody remembers to configure it.

Financial data withheld by design

Budgets, expenses, timesheets and billable rates are granted only to the two approver roles, and are omitted entirely from the product manager, quality lead and the stakeholder seat handed to external clients.

Approval as a gated, recorded action

Budget approval, timesheet approval and quality-gate review are privileged actions distinct from editing, and each records who signed off and when on the record itself.

Single sign-on and directory sync

SAML single sign-on and SCIM provisioning are inherited from the platform, so joiners and leavers are handled by your directory rather than by a manual member list.

Attributable audit trail

Actor identity travels with every request and is used for audit and owner-scoped filtering, so changes are attributable and a history exists for anything that moved.

Data minimisation and portability

Only the data needed to run the service is collected, it is never sold, and export and deletion paths exist so a workspace's contents can be taken out as readily as they were brought in.

09 Stated plainly

PlanMagnet holds no security certification today. It is designed toward controls of a SOC 2 style and built to align with European data-protection principles, and a maintained subprocessor list is published, but neither SOC 2 nor ISO 27001 has been attested, and incumbents in this category already hold both. That is a genuine gap for a regulated buyer and should be weighed as one. Three further limitations are worth stating rather than discovering: time entries and timesheets are one governed resource, so a contributor able to log their own hours can also read colleagues' hours and billable rates; custom field definitions and their per-record values are likewise one resource, so populating a field requires the ability to define one; and around two dozen destructive tools reachable by an external agent do not yet require a confirmation step, so an agent acting inside its granted rights can delete faster than a person would. Each of these is a known constraint with a specified fix, not an oversight.

10 One product, many front doors

PlanMagnet is one product with several front doors. The web workspace is where the work happens; everything else exists so the product can be reached from wherever a person or an agent already is. Statuses below are honest — several of these are planned for launch rather than downloadable today.

Web workspace

The full product in the browser: roadmaps, backlogs, boards, sprints, releases, portfolios, risks, budgets, timesheets and dashboards. This is the primary surface and the one the launch is built around.

Federated GraphQL API

The published contract every other surface is built on, with scoped keys and the same permission enforcement the app gets. Live today.

Agent tool server

A published Model Context Protocol server exposing 155 typed planning tools across 18 areas of the product, plus the shared platform tools, so an external agent can complete a real task rather than stall at the first cross-cutting concern.

Command-line tool

Drive the workspace from a terminal or a pipeline with the same authorisation the apps use. Early and not yet ready for production reliance.

Node SDK

A typed, promise-based client for any Node or Bun project, for internal tooling and automation over the same API. Early.

Python SDK

An async client for data pipelines, delivery analysis and notebooks. Early, and targeting recent Python versions.

Editor extension

A planned sidebar for the code editor: assigned issues for the current sprint, a jump to the parent epic, status updates and commit-to-item linking without leaving the editor.

Browser extension

A planned side panel to check the current sprint, log a quick issue or tag a page for triage from any tab.

Desktop companion

A planned native application mirroring boards with local notifications and offline tolerance, so a dropped connection does not stall the team mid-sprint.

Mobile applications

Planned native applications for reviewing a sprint, approving a release and capturing an idea into the backlog away from a desk. The launch is web-first and says so.

Documentation

Product and API documentation published alongside the product, including the machine-readable formats an AI assistant can consume directly.

11 Plans

Billing is per editor: only people who create or change work are charged, and viewers are free on every tier. These are launch list prices in the primary currency band and are planning figures ahead of launch.

FREE

\$0

Small teams proving the model, and any org that wants a governed workspace before it wants a contract.

- Up to 10 editors
- 3 projects, unlimited work items
- 2 boards, 1 sprint, 1 roadmap
- Role-based access, single sign-on and audit inherited from the platform
- Unlimited free viewers
- Import from Jira, Linear, Asana, GitHub and Trello

PRO

\$9/editor/mo · \$7.50 annual

A single team running real delivery: product and engineering on one backlog with a roadmap above it.

- Unlimited projects
- Scrum and Kanban boards with sprints and story points
- Roadmaps linked to the work that delivers them
- Custom fields and custom workflows
- Time tracking
- Automation rules, metered
- Unlimited free viewers

BUSINESS

\$16/editor/mo · \$13 annual

Multi-team delivery, agencies and programme offices that need governance and finance in the same workspace.

- Portfolios and programmes
- Cross-project dependencies
- Risk register and quality gates
- Budgets, expenses and timesheet approval
- Dashboards and delivery analytics
- Single sign-on sandbox
- 99.9% availability target

ENTERPRISE

From \$26/editor/mo list · custom annual

Organisations with governance, residency or procurement requirements that shared infrastructure cannot satisfy.

- Unlimited everything
- Advanced governance and audit
- Data residency options
- Dedicated tenant or self-host options
- Directory provisioning at scale
- Negotiated support terms

11 How billing works

Per-editor billing, free viewers

Only people who create or modify work are billed. Stakeholders, clients, executives and commenters have read and comment access at no cost on every tier, which is where most of the cost difference against seat-based tools comes from.

Four regional price bands

Pricing is banded by region, including native rupee pricing for India, so the product is not priced as though every market were the same one. The figures above are the primary band.

Planning figures, honestly labelled

PlanMagnet is pre-launch. These are launch list prices from the current pricing plan and may change before general availability based on cohort data and finance confirmation.

What sits outside the plan

Marketplace purchases, add-ons and any usage over a plan's metered allowances — automation runs in particular — are billed separately. Enterprise pricing is custom and depends on compliance, residency and support requirements.

12 Why teams choose it

Product and engineering share one data model

A feedback inbox with voting and priority scoring, roadmaps and versioned releases sit beside boards, sprints, issues and dependencies in one workspace with one identity and one audit trail. The usual alternative is an issue tracker plus a wiki plus a roadmap tool, kept in step by hand.

Delivery operations are included, not sold separately

Portfolios, programmes, a risk register, quality gates, budgets and expenses, time tracking and timesheets, custom fields, templates and automation rules all ship as real parts of the workspace. Dedicated planning tools in this category typically do not carry that operational depth natively.

Migration ships as a product feature

Import and export both work for Jira, Linear, Asana, GitHub and Trello. Most evaluations in this category stall on migration pain; answering it in-product on day one makes adopting PlanMagnet a reversible experiment rather than a commitment.

Governance from the free tier

Tenant isolation, role-based access on 21 resource types, single sign-on, directory sync, audit and quotas are inherited from the platform underneath rather than being an enterprise upsell, so the free tier is already a governed workspace.

Viewers cost nothing, on every tier

Executives, clients and commenters read, follow and vote without a seat. For teams with many stakeholders per editor, this is the difference between a plan you can extend to the whole org and one you ration.

AI that cannot write to your plan

The live AI capability has no code path that creates a work item, and every call writes an append-only record before returning. This is enforced by structure rather than promised by policy, which is an unusual starting position in this category.

AI grounded in your own delivery history

Because the roadmap, sprints, dependency graph and velocity are the same records, a planning suggestion can point at the work behind it. That is a structural advantage over assistants bolted onto trackers whose plan lives in another system.

Regional pricing, not a single dollar list

Four price bands including native rupee pricing make the product viable in markets where dollar-led competitors are priced out of reach.

12 Honest answers to hard questions

You are pre-launch. Why would we bet on you?

You should weigh that honestly, and we are not going to dress it up: there is no proven scale, no certification and no public track record. What there is, is a working delivery workspace with import and export from the tool you are on now, so an evaluation costs you a real backlog import and nothing else, and reversing it is a supported operation rather than a rescue project.

Your competitors have SOC 2 and ISO 27001. You do not.

Correct, and for a regulated buyer that may be decisive today. Controls are being designed toward a SOC 2 shape and the product aligns with European data-protection principles, with a published subprocessor list — but neither certification has been attested, and we will not imply otherwise. If certification is a hard gate for you, we are not yet the right answer.

Most of your AI is marked planned.

Three of the four committed decisions are indeed not built, and the app has no assistant. What is built is the part most vendors skip: a single governed call path, strict output validation, an append-only record of every call including failures, and a live capability that structurally cannot write to your backlog. We would rather ship the governance first and the features second.

No spec editor and no objectives framework. Those are table stakes for us.

Both are specified and committed, and neither exists in the product today, so if your evaluation turns on them you should discount PlanMagnet accordingly. What does exist is the link from an initiative to the work delivering it, which is the connection those features usually exist to provide.

We would be trading a mature ecosystem for a young one.

Yes. The largest incumbent has thousands of marketplace applications and we will not match that breadth at launch. The counter-argument is that a large part of that marketplace exists to reconnect products that were sold separately in the first place — roadmaps to tickets, docs to work, feedback to plan — and those connections are native here.

Mobile matters to our field and executive users.

The launch is web-first. Native mobile applications are planned, not shipped, and competitors in this category ship mature ones today. The web workspace is responsive, but if a native mobile experience is a requirement rather than a preference, that gap is real now.

Our engineers will reject anything slower than what they use today.

A fair test, and one we have to pass rather than argue. The product is opinionated about boards, sprints and keyboard-driven work for exactly this reason, but performance under a real backlog is something to verify in an evaluation rather than take on assurance.

Can an AI agent do damage if we connect one?

It cannot exceed the rights of the account you connect it with, because every tool call is permission-checked exactly as a person's would be. The honest gap is inside those rights: around two dozen destructive tools have no confirmation step or dry-run mode yet. Until that lands, scope the account you hand an agent narrowly, and treat that as the control.

Where it sits against the alternatives

VS ENTERPRISE AGILE TRACKERS

They are the benchmark for agile execution and hold the compliance credentials, but a full picture usually means the tracker plus a wiki plus a discovery product plus an advanced roadmap add-on, and real cost lands well above list. PlanMagnet puts those in one workspace with one bill and imports your existing issues to start. The concession is honest: their marketplace breadth, certification and proven scale are all ahead of ours.

VS MINIMAL ENGINEERING ISSUE TRACKERS

They are fast, polished and deservedly loved, and they deliberately leave out product-management depth. PlanMagnet adds the feedback inbox, versioned releases and delivery-operations layers they skip, and imports directly from them. The concession: their roadmap already rolls progress up from issue completion, which PlanMagnet stores but does not yet compute, and we do not claim roadmap superiority over them.

VS PRODUCT-STRATEGY AND ROADMAP SUITES

They set the benchmark for roadmap presentation and strategy hierarchy, but engineering execution is a separate add-on and the assembled stack gets expensive quickly. Boards, sprints, issues and releases are native here, at one price, with free viewers on every tier. The concession: their roadmap presentation depth and dedicated research modules are genuinely better than ours.

VS FEEDBACK-TO-ROADMAP PLATFORMS

They are excellent at turning customer signal into a prioritised plan, and then deliberately push execution to a separate tracker.

PlanMagnet keeps the feedback inbox next to the backlog it feeds, so a request can be followed through to the sprint that ships it.

The concession: their feedback aggregation, public portal and AI over that signal are more mature than what we ship.

VS GENERAL-PURPOSE WORK PLATFORMS

They are extremely flexible across departments, and that flexibility is exactly why software delivery needs a separate developer product bolted on, with viewers often billed as seats. PlanMagnet is opinionated about the shape of product and engineering work instead. The concession: their cross-functional breadth, automation maturity and mobile applications are all ahead of ours.

VS EVERYTHING-APP WORK PLATFORMS

They pack tasks, docs, whiteboards, chat and sprints into one aggressively priced product, which is a real advantage on breadth per pound. PlanMagnet trades breadth for an opinionated delivery model and shipped delivery-operations depth. The concession: they offer more native modules at a lower price, and a far larger template and automation library today.

START TODAY

Start with your real backlog, not a demo dataset

PlanMagnet is pre-launch and taking waitlist registrations. The fastest way to judge it is to bring the work you already have.

- Join the waitlist and tell us which tool you are on today
- Import a real project from Jira, Linear, Asana, GitHub or Trello
- Run one sprint end to end, then read the velocity it produced
- Hand a stakeholder a free viewer seat and see whether the roadmap answers their question

WAITLIST

planmagnet.com/waitlist

SCENARIOS IN FULL

planmagnet.com/use-cases

AI, WITH STATUSES

planmagnet.com/ai

SECURITY POSTURE

planmagnet.com/security

DOCUMENTATION

docs.planmagnet.com
