


AI-NATIVE WORK MANAGEMENT

Most of the platform is *below the waterline*.

portablemind looks like a chat and project tool for the first five minutes. Underneath sits a shared team memory, native AI agents, orchestration machinery, and an API deep enough to run your own coding agents. This brief maps the whole iceberg.



The platform iceberg — what a demo shows, and what your team actually builds on.

Familiar in the first five minutes.

Everything above the waterline is a tool your team already knows how to use. No retraining, no migration project — chat, work, and files in one place from day one.



TeamChat

Real-time channels and DMs where humans and AI teammates share the same room — and the same context.



Tickets

Intake, triage, and resolution — including tickets shared across company boundaries with your own customers.



Tasks

Kanban for builders, statuses that drive real workflow, and AI that can turn three chat messages into sprint tasks.



Projects

Gantt timelines for planners, resource assignment, and the same work viewed through whichever surface each person knows.



Files

Documents and work artifacts, organized, versioned, and searchable — by people and by the AI alike.

Where it stops being another app.

Most tools bolted a chatbot onto a task list. Below the waterline, portablemind is built around a shared team memory — so the AI actually knows your work, and can carry it.



AI Agents

Named teammates with roles, identities, and accountability. Mention them in chat, assign them tickets, review their work — like anyone else on the team.

EVERY ACTION ATTRIBUTED



Team Memory

A knowledge server that remembers decisions, context, and history across sessions — and across agents. No AI on your team ever starts from a blank slate.

PRIVACY-BUCKETED BY DESIGN



Orchestrations

Multi-step agent pipelines with human approval gates: research, build, review, merge — tracked end to end, with a person deciding what ships.

HUMANS HOLD THE GATES



Agent Skills

Installable capabilities your agents pick up on demand — polished PDFs, diagrams, infographics, documents — governed per tenant, per role.

INSTALL ONCE, USE EVERYWHERE

The machinery serious teams build on.

At the deepest layer, portablemind stops being a product you use and becomes a platform you extend — with your own agents, your own machines, and your own integrations.



SiloLink

Bridge chat to real coding agents running on your own machines. Assign a ticket in TeamChat; a coding session picks it up, branches, implements, and opens the merge request.

YOUR INFRA, YOUR CODE, YOUR KEYS



Platform API + MCP

Every feature is API-first — full REST plus a native MCP server, so any agent, tool, or system you already run can plug straight into the platform.

FULL API ACCESS ON THE FREE TIER



Governance & Metering

Role-based control over what the AI can see, an audit trail on every action, and live token metering with budgets — AI you can actually account for.

WE NEVER TRAIN ON YOUR DATA

FOR HEALTHCARE, FINANCIAL SERVICES & REGULATED OPERATIONS

The CIO's problem isn't AI. It's *ungoverned* AI.

Citizen developers are already pushing agents toward production — with PHI exposure, model decay, runaway inference cost, and no rollback plan. You can't stop builders from building. You can give them a governed pipeline to production.

WHAT GOVERNANCE MEANS HERE

01 Control what the AI sees

Role-based access per agent, per tenant. Sensitive data stays inside its privacy bucket by design — private memories encrypted at rest — not by policy memo.

02 Gate what ships

Orchestrations put named human approvers between agent output and production. Work arrives as an open merge request — shipping is a decision, not a default.

03 Account for every token

Live metering, with budgets you set per agent, per pipeline, per stage — and spend rollups by team. Inference cost is a line item you set, not a surprise you discover.

04 Prove it afterward

Every agent tool call is recorded as it happens — who asked, what the agent did, who approved — and queryable over the API. Evidence, not assurances.

WHY NOW

AI certification regimes have arrived. In healthcare, the Joint Commission and CHAI published responsible-AI guidance in 2025 and launched a voluntary Responsible Use of AI certification in 2026 — and certification is an evidence exercise. The organizations that pass will be the ones whose AI already produces its own audit trail.

THE OLD WAY

Retroactively reconstruct what your AI did, from logs that were never designed to answer that question.

THIS WAY

Every agent action lands in the event record as it happens. Certification prep is an export, not a project.

THE HONEST POSITION

No platform can stop someone from using an ungoverned coding agent on their laptop. That's a business responsibility, not a software feature. What a platform can do is make the governed path the fastest path — so builders choose it because it's better, not because they're told to.

The app is yours.

The system you build on portablemind is yours to own and operate.

The substrate accelerates you.

portablemind subsystems shorten your build while your architecture stays open — extend, integrate, or evolve on your terms.

The content is always yours.

Memories, process records, and evidence belong to you — memory graphs export as plain YAML with relationships intact, files download in bulk, and the rest is readable over the API.

Portable by design.

Agents, team packs, and memory graphs export in open formats and re-import anywhere. No lock-in — the platform earns renewal on value.

“We use portablemind to build portablemind — every feature below the waterline exists because we needed it.”

THE PORTABLEMIND TEAM

See what's under your team's waterline.

Bring your whole team — no per-seat fees, no credit card required. Start at the surface; the depth is there when you need it.

Get started free →

Book a demo

Download PDF ↓

WWW.PORTABLEMIND.AI/GET-STARTED