Roman Manchenko

AI Executive — Head of AI / VP of AI | Agentic Platforms, AI Products & Enterprise AI Adoption

+1 (646) 708-2890 | manchenko.roman@gmail.com | Berkeley Heights, NJ (NYC area)

LinkedIn: https://www.linkedin.com/in/roman-manchenko-2a830030/

SUMMARY

Engineering executive, product thinker, and founding team member with 20 years in software and 10+ years leading engineering at scale, including organizations of 70+ people across 20 countries. As CTO of RebelMouse, a content platform for digital publishers and enterprise brands, owns AI end to end: the strategy the CEO approved, the agentic platform (agents, context and retrieval, evals, permissions, feedback loops), the products and internal applications built on it, their pricing, and the change management that got people to actually use them. Takes AI from impressive demo to production system, ties it to business outcomes, and still ships production code daily.

CORE EXPERTISE

AI Strategy & Value

- Multi-year AI strategy; build / partner / buy decisions; AI investment tied to measurable outcomes
- AI product pricing and packaging; executive, client, and prospect advisory

AI Platform & Architecture

- Agents and orchestration: Claude Agent SDK, MCP, skills, human-in-the-loop approvals
- Context and retrieval (embeddings, RAG, enrichment); evals (golden sets, LLM-as-judge, per-version metrics)
- Permissions, governance, and feedback systems; multi-provider model routing and usage metering

AI Products & Adoption

- AI-native products built on proprietary content and data; moving from demo to production
- Org-wide change management; outcome metrics; workflow redesign across support, delivery, content, engineering

Engineering Leadership & Foundation

- Distributed international teams; managing managers; hiring and succession; SOC 2, 99.99% uptime
- TypeScript, Node.js, Python, Go · AWS, EKS, Terraform · PostgreSQL, MongoDB, Redis · React, Astro

PROFESSIONAL EXPERIENCE

Chief Technology Officer, RebelMouse Inc | New York, NY

03/2018 – PRESENT

Content platform and CMS for digital publishers and enterprise brands. Reports to the CEO; owns engineering, DevOps, QA, product support, and company AI strategy.

AI Strategy & Business Outcomes

- Authored the company's AI-native strategy: repositioned it from a commodity CMS to an agentic platform where AI agents originate the work and people make the final call. Pitched it to the CEO, secured approval, and now own execution end to end — architecture, roadmap, pricing, and prospect demos.
- Made the build-vs-buy calls: built on the Claude Agent SDK and MCP instead of a proprietary orchestration stack; treated the workflow engine as a commodity and invested where competitors can't copy — reusable playbooks with measured quality and an agent-ready content platform.
- Designed pricing and packaging for the AI platform: a credit model that meters LLM tokens and compute, per-model cost, and a bring-your-own-key enterprise tier that lets clients run model spend on their own Azure OpenAI or AWS Bedrock contracts.
- Re-engineered client delivery as deal sizes compressed: onboarding and site builds move to agent-run workflows with people reviewing and approving — design-fidelity QA against Figma, WordPress content migration, launch checklists — so smaller contracts stay profitable.
- Restructured from a 70+ person organization to a lean, higher-leverage team as AI productivity reduced the headcount needed to ship.

Agentic Platform — Architecture & Production

- Architected and shipped the company's agentic operating system — workflow engine, agents, skills, MCP tools, task inbox, and approval gates — used by internal teams and exposed to clients inside the CMS.

- Introduced "playbooks": an interview-style agent captures how a process really works, and a generator compiles it into a workflow where every step has a human or agent owner — automation grows one step at a time.
- Built the context and retrieval layer: MCP servers that expose content and CMS operations to agents, an embeddings service with collections for retrieval, and an event-driven enrichment layer that stores LLM output on content records instead of recomputing it on every run.
- Closed the evals and feedback loop in production: in-product thumbs-up/down and notes become outcome labels on each run; one click turns a failure into a golden-set case with an auto-derived rubric; LLM-as-judge replays score new versions against the golden set.
- Set permissions and governance for agents: tiered skills (platform / organization / site / user) with explicit precedence, per-user identity on agent runs and on any tool that touches customer data (no shared tokens), and approval gates on write actions.

AI Applications in Production

- **Support triage agent:** every new Zendesk ticket is triaged, checked for duplicates across the knowledge base, Jira, and past tickets, assessed as a product bug or not, then filed in Jira with a Slack alert — with a human gate when information is missing. Hardened through 40 production versions.
- **CMS Copilot:** a context-aware assistant in the customer dashboard that explains, navigates, and acts ("show me my drafts" applies the real filters). Capabilities climb a ladder — explain → navigate → read → write — and writes require in-chat confirmation under the user's own permissions. Three user-facing modes: Classic, Copilot, Autopilot.
- **Personal assistants:** any employee can build "just for me" workflows that run under their own identity and permissions, starting from a guided "first workflow" experience.
- **AI-ready content delivery:** llms.txt and per-page Markdown for AI crawlers and answer engines, plus a public MCP server and OpenAPI/SDK so publishers' content and the platform itself are operable by agents.

AI Adoption & Operating Model

- Moved the company from "AI as a faster chatbot" to agentic workflows as the default way work gets done. Diagnosed adoption resistance as identity-driven, not output-driven, and redesigned the operating model around outcomes: every initiative tied to a result, tool choice left to engineers (Claude Code, Codex, Gemini), AI-adjusted estimation, and an outcome-based output dashboard with a 3–5x greenfield / 2–3x legacy bar.
- Lead by example: authored 151 of the team's 774 merged pull requests across 31 repositories in the last 30 days — about 180 merges a week from a small product-engineering team.
- Took adoption beyond engineering: weekly show-and-tell workshops, an idea board, and a hub inside Google Workspace (where staff already work), tracking workflows shipped per week and unique contributors.
- Treat AI features as products — measure, and cut what doesn't earn usage: stopped a dashboard-by-dashboard copilot rollout in favor of shared MCP tools and skills that every surface reuses.

Architecture & Platform Modernization

- Designed an AI-native rendering engine: AI emits structured JSON page definitions rendered by a deterministic, typed component library (Astro) — engineered for >90% AI build success versus ~60% with freeform generation.

Team Leadership & Operations

- Built the engineering organization to a peak of 70+ people across 20 countries; today lead 15 across 9 countries — software engineering, DevOps, QA, and product support.
- Maintain 99.99% uptime; led the migration from monolith to AWS EKS microservices, CI/CD, and SOC 2 certification.
- Cut time-to-hire by 40%; partner with sales on enterprise deals, discovery calls, and tailored AI demos for prospects.

Earlier Roles

- **VP of Engineering**, RebelMouse (12/2016 – 03/2018) — led a cross-functional team of 15.
- **Managing Director, Engineering**, RebelMouse (03/2014 – 12/2016) — took features from idea to production.
- **Frontend Team Lead**, RebelMouse (10/2011 – 03/2014) — founding team member; built the first product.
- **Senior Frontend Engineer**, The Huffington Post (10/2010 – 10/2011) — architected newsGlide, HuffPost's SPA.
- **Full-stack Web Engineer**, Bertelsmann (08/2006 – 10/2010) — first engineer; grew to lead the team.

EDUCATION

M.S. Computer Science (2005) and **B.S. Mathematics & Management** (2004) — Kharkiv National University of Radio Electronics