

AI Needs Radar

Weekly insights from global tech communities

2026.05.17 — 2026.05.24

Issue #4 • Generated 2026-05-24

619

discussions analyzed

26

needs gaps surfaced

6

categories covered

6

big-company opportunities

// TOP 10

Top 10 Needs by Heat

The hottest needs across all categories this period.

#01	Multiple AI agents can't collaborate	9/10
AI Development	UP • Using Claude Code, Codex, and Aider together causes conflict	
#02	No security boundary for sensitive agent actions	8/10
AI Development	UP • Agents execute commands, send emails, and modify databases s	
#03	No reliable fallback when AI agent errors	8/10
AI Development	UP • Agents get stuck in error loops, wasting tokens and time unt	
#04	Can't see which skills and tokens agent uses	7/10
AI Development	UP • Skills reload every session, wasting context without knowing	
#05	Code search grep is too primitive for agents	7/10
AI Development	UP • Agent grep reads entire code, exploding tokens and pulling i	
#06	Year of video footage lost on hard drive	7/10
AI Content	UP • Thousands of hours of video recorded, but finding a specific	
#07	Colleagues waste time with AI-answerable questions	6/10
AI Tools	UP • Team members ask humans directly instead of querying AI fir	
#08	AI-generated internal tools are low quality and dangerous	6/10
AI Development	UP • AI writes an ops panel easily but might directly operate the	
#09	AI agent reloads skills wasting context	6/10
AI Development	UP • Each user command reloads the same skills, causing insuffici	
#10	Unsure which grad school to choose	6/10
AI Education	UP • Want to advance education but school info is scattered and m	

// AI TOOLS

AI Tools

4 needs • sorted by heat

1. Colleagues waste time with AI-answerable questions • Heat 6/10

Team members ask humans before AI for trivial questions

- PAIN** Interrupted multiple times daily to answer questions AI could solve instantly
- EXISTING** LMGTFY useless, no AI version, suffer in silence
- OPPORTUNITY** Integrate with Slack/WeCom to auto-detect AI-answerable questions and push AI response first
- COMPETITORS** Slack AI won't do this, Copilot Chat only personal use

For: Small Team • Difficulty 3/5 • UP

2. Auto-generate step-by-step docs from screen recordings • Heat 5/10

Recorded videos still need manual screenshots and SOP writing

- PAIN** Existing AI SOP generators only write 'click here' without business context
- EXISTING** DIY video2docs projects, no mature product on market
- OPPORTUNITY** Record once, auto-add captions, arrows, and exception branches
- COMPETITORS** Scribehow, Tango only do screenshots, no video intelligence

For: Small Team • Difficulty 4/5 • UP

3. Notes scattered across four apps, impossible to find • Heat 5/10

Notion, Obsidian, WeChat favorites scattered, can't find half-year-old ideas

- PAIN** Keywords fail after 8 variations, timeline completely messy
- EXISTING** Notion AI only queries current workspace, Obsidian needs manual backlinks
- OPPORTUNITY** Unify all note sources, return original location and similar content for any query
- COMPETITORS** Mem, Reflect only bind partial sources, become silos after import

For: Small Team • Difficulty 4/5 • UP

4. Distracted by editor while trying to focus on writing • Heat 3/10

Every computer session becomes fixing blog toolchain instead of writing

PAIN Wasting writing flow on CSS tweaks and plugin changes, paper notebooks inconvenient

EXISTING Markdown editor + strict discipline, but mind can't control hands

OPPORTUNITY Pure writing interface with only text rendering, no settings menu

COMPETITORS iA Writer, Ulysses still allow font style adjustments, not minimal enough

For: Solo • Difficulty 2/5 • STABLE

// AI DEVELOPMENT

AI Development

9 needs • sorted by heat

1. Multiple AI agents can't collaborate • Heat 9/10

Using Claude Code, Codex, Aider simultaneously, they overwrite each other in the same project

- PAIN** No unified orchestrator, A changes then B reverts, manual sync is exhausting
- EXISTING** Ait project does this but is very early
- OPPORTUNITY** Provide a control plane for multiple agents to work in parallel on different worktrees
- COMPETITORS** Ait is local CLI only, no cross-process agent collaboration

For: Enterprise • Difficulty 5/5 • UP

2. No security boundary for sensitive agent actions • Heat 8/10

Agents can execute commands, send emails, modify DBs, relying solely on prompt constraints

- PAIN** System prompt is not a security boundary, one prompt injection can delete the database
- EXISTING** Enforra, Capframe just open-sourced, no adoption yet
- OPPORTUNITY** Policy check before tool execution, like Android permission dialogs
- COMPETITORS** Ollama, vLLM have no built-in security layer, LangGraph is manual

For: Enterprise • Difficulty 4/5 • UP

3. No reliable fallback when AI agent errors • Heat 8/10

Agent loops on error branches, wasting tokens and time

- PAIN** No state machine control, sub-agent stuck without returning to main flow, only escape to interrupt
- EXISTING** Statewright uses state machine constraints but requires manual modeling
- OPPORTUNITY** Auto-generate state machine from prompt, degrade to predefined paths on error
- COMPETITORS** LangGraph too flexible with no enforced boundaries, Claude Code has no fallback

For: Enterprise • Difficulty 5/5 • UP

4. Can't see which skills and tokens agent uses • Heat 7/10

Skills reload every session, unknown which skills are unused

- PAIN** No dashboard showing skill usage, token cost, success/failure rates
- EXISTING** Skilled CLI can query sessions but only own sessions
- OPPORTUNITY** Provide team-level agent usage analytics panel to optimize skill configuration
- COMPETITORS** Semble only does code search, Skilled has no cloud

For: Small Team • Difficulty 3/5 • UP

5. Code search grep is too primitive for agents • Heat 7/10

Agent grep reads entire code blocks causing token explosion, often pulls irrelevant files

- PAIN** Grep returns 100 lines of matches, agent can't tell which is the actual function
- EXISTING** Semble compresses 98% tokens with code structure, but initial version doesn't support complex queries
- OPPORTUNITY** Build code semantic search engine returning function signatures and call relationships
- COMPETITORS** Sourcegraph Cody is IDE search, doesn't support direct agent calls

For: Small Team • Difficulty 4/5 • UP

6. AI-generated internal tools are low quality and dangerous • Heat 6/10

Writing an ops dashboard with AI is easy, but it might directly operate production databases

- PAIN** AI-written internal tools have no access control, one mistake wipes data
- EXISTING** Manual permissions plus code review, but agent-written code is rarely reviewed
- OPPORTUNITY** Provide a read-only data validation layer, AI can only read unless specially authorized
- COMPETITORS** Retool allows customization but no AI security layer, Airplane has no similar protection

For: Enterprise • Difficulty 4/5 • UP

7. AI agent reloads skills wasting context • Heat 6/10

Every user instruction reloads the same skills, context window insufficient

- PAIN** After compression, reloading skills, half the context is duplicate content
- EXISTING** Manually trim skill descriptions, but agent memory is insufficient
- OPPORTUNITY** Skill caching mechanism: load once per session, auto-restore after compression
- COMPETITORS** Claude Code, Codex don't deduplicate skill loading

For: Small Team • Difficulty 3/5 • UP

8. Fragmented lock file support across package managers • Heat 5/10

Bun's lock file is ignored by some tools, causing CI inconsistency

- PAIN** Each package manager has its own lock file format, tools must manually support 5 formats
- EXISTING** Maintainers hardcode each lock file, adding a new one requires source changes
- OPPORTUNITY** Unified lock file parsing interface, tools only depend on this interface
- COMPETITORS** Pnpm doesn't support abstraction, npm has no such abstraction

For: Small Team • Difficulty 3/5 • STABLE

9. Cross-platform CLI lacks unified DDC command support • Heat 4/10

Want to control monitor KVM on Windows, but ddcutil only works on Mac/Linux

- PAIN** When automating dual-machine shared monitor, Windows has no equivalent command-line tool
- EXISTING** M1ddc only for Mac, Windows requires GUI software or manual switching
- OPPORTUNITY** A cross-three-platform DDC CLI supporting scripted input source switching
- COMPETITORS** Ddcutil only Linux, monitorCTRL only Windows GUI

For: Solo • Difficulty 3/5 • UP

// AI HARDWARE

AI Hardware

3 needs • sorted by heat

1. Smart home switch status falsely shows success • Heat 5/10

Mi Home shows light off, but it's still on

- PAIN** Command status and actual device status out of sync, failure reported later
- EXISTING** Wait seconds and check logs? No real-time status feedback
- OPPORTUNITY** Device status uses two-step confirmation, UI updates only after receiving acknowledgment
- COMPETITORS** Mi Home and HomeKit use optimistic updates without waiting for device response

For: Enterprise • Difficulty 3/5 • DOWN

2. Apple Watch battery can't last long hike • Heat 4/10

Sports mode drops below 20% in 6 hours, needs mid-trip charge

- PAIN** Anxious about running out of battery during mountain hiking tracking
- EXISTING** Carry power bank and cable, but inconvenient outdoors
- OPPORTUNITY** More power-efficient sports mode or dual battery design
- COMPETITORS** Garmin lasts weeks but ecosystem weaker than Apple

For: Enterprise • Difficulty 4/5 • UP

3. Wi-Fi router reconnection triggers light strip • Heat 3/10

Appliances keep pre-power-off state, but Mi light strip always turns on

- PAIN** Gateway reconnection triggers device auto-on, can't find controller
- EXISTING** Replace official power supply? Not sure if effective
- OPPORTUNITY** Provide 'force off on power-up' low-level mode not overridable by automation
- COMPETITORS** Mi Home lacks this option, HomeKit has no forced default state

For: Small Team • Difficulty 2/5 • STABLE

// AI CONTENT

AI Content

4 needs • sorted by heat

1. Year of video footage lost on hard drive • Heat 7/10

Thousands of hours recorded, finding a scene relies on memory

PAIN No local private index, uploading to cloud risks leaks, only manual search

EXISTING Some use Gemma4+frame extraction to index, but most can't

OPPORTUNITY A local video content indexer supporting natural language search

COMPETITORS Google Photos only searches by time/object, can't understand events

For: Solo • Difficulty 4/5 • UP

2. Video training for volunteers but VPS can't handle • Heat 5/10

Training videos on Oracle VPS, bandwidth insufficient with many viewers

PAIN Large video files overwhelm server performance

EXISTING Bilibili or Youku but no progress tracking, self-hosted streaming complex

OPPORTUNITY Lightweight video hosting with playback progress tracking, API embeddable

COMPETITORS Vimeo OTT too expensive, Cloudflare Stream more professional but no progress management

For: Small Team • Difficulty 3/5 • UP

3. Content scraped by AI, no traffic back • Heat 5/10

Hard-written tech articles, AI summarizes them, no one visits my site

PAIN Google/Bing use content to train AI, no referral traffic to authors

EXISTING Add login wall or robots.txt, but inconvenient for own access

OPPORTUNITY Content watermarking + access stats + AI output attribution, forming paid citation mechanism

COMPETITORS Cloudflare Turnstile can verify but can't track citations

For: Small Team • Difficulty 4/5 • UP

4. AI generates hollow content, wasting creativity • Heat 4/10

Want AI to help write story outlines, but output is bland templates

PAIN Tuned prompts for hours still feel plastic, tweaking wastes time

EXISTING Manually polish prompts, but need to re-tune for each project

OPPORTUNITY Visual outline editor, let AI fill in details instead of generating full text

COMPETITORS Notion AI, Copy.ai generate finished products directly, lack co-creation feel

For: Solo • Difficulty 3/5 • STABLE

// AI LIFESTYLE

AI Lifestyle

3 needs • sorted by heat

1. Insurance AI gives product names but no selection help • Heat 5/10

Asking AI for insurance yields product names without clear comparison.

PAIN Complex terms, AI can't analyze deductibles, renewal conditions, or pre-existing conditions.

EXISTING Alipay Good Health page compares, but different companies are separate.

OPPORTUNITY Input age and scenario, AI compares clause differences and gives risk warnings.

COMPETITORS Simple insurance platforms lack clause comparison; Huize brokers are pricier.

For: Small Team • Difficulty 3/5 • UP

2. Zero-based men's skincare beginners can't find first step • Heat 4/10

Online skincare tutorials are too complex; don't know what to buy or order.

PAIN Information bubble full of ads or too technical, afraid to buy wrong.

EXISTING Asking AI gives generic steps but lacks personalized recommendations.

OPPORTUNITY Photo face analysis + budget + skin type, directly give purchase list and usage.

COMPETITORS Beauty Evolution only has ingredients, no personal customization.

For: Solo • Difficulty 2/5 • NEW

3. Friend borrows money without proof, hard to collect later • Heat 3/10

Lent money to acquaintance without proper IOU, delays cause stress.

PAIN After-the-fact IOUs have weak legal effect; e-signatures feel troublesome.

EXISTING Tencent e-sign mini-program, but many don't know or find steps too many.

OPPORTUNITY One-click generate legal e-IOU, auto-remind repayment after signing.

COMPETITORS Tencent e-sign and Jiedaibao are cumbersome, no auto-reminders.

For: Solo • Difficulty 2/5 • STABLE

// AI EDUCATION

AI Education

3 needs • sorted by heat

1. Unsure which grad school to choose • Heat 6/10

Want to advance degree but school info is scattered

- PAIN** Admission guides are lengthy, hard to match background and goals
- EXISTING** Forums and Zhihu posts, but outdated and incomplete
- OPPORTUNITY** Input field and location, AI recommends best schools and generates timeline
- COMPETITORS** Kaoyanbang and Yanzhaowang have data but no smart matching, cross-region analysis weak

For: Small Team • Difficulty 3/5 • UP

2. AI teaches outdated Rust syntax • Heat 4/10

AI-generated Rust code uses old patterns, hard to detect

- PAIN** AI suggests idioms that may be deprecated without warning
- EXISTING** Manually check Rust edition changelog
- OPPORTUNITY** AI tutor auto-annotates version compatibility and points to latest practices
- COMPETITORS** Copilot Chat lacks version awareness, Rustlings is manual exercises

For: Small Team • Difficulty 4/5 • NEW

3. No immersive practice for new language • Heat 3/10

Stuck after tutorials, lack real projects to apply skills

- PAIN** Errors only appear in production, never encountered during learning
- EXISTING** LeetCode and toy projects, but too simple
- OPPORTUNITY** Generate runnable project bugs for students to fix and learn
- COMPETITORS** GitHub Copilot writes code directly, doesn't teach debugging

For: Small Team • Difficulty 3/5 • STABLE

// BIG PICTURE

Big Company View (Enterprise)

Directions that need strategic synergy, scale and brand — beyond what solo or small teams can pursue.

1. Multiple AI agents can't collaborate • Heat 9/10 • AI Development

Using Claude Code, Codex, Aider simultaneously, they overwrite each other in the same project

STRATEGIC Collaboration layer for AI coding, team adoption decision point, TAM covers all software teams

INCUMBENTS Cursor, Copilot are single-agent, can't parallelize tasks

TIMING Early agent explosion, no standard collaboration protocol yet

2. No security boundary for sensitive agent actions • Heat 8/10 • AI Development

Agents can execute commands, send emails, modify DBs, relying solely on prompt constraints

STRATEGIC Every agent deployment needs this, moat in policy engine and compliance audit

INCUMBENTS Semgrep has no runtime policy, Copilot has no call blocking

TIMING More agents entering production, security incidents rising

3. No reliable fallback when AI agent errors • Heat 8/10 • AI Development

Agent loops on error branches, wasting tokens and time

STRATEGIC Determines if agent can be used for production-critical tasks, large TAM and high stickiness

INCUMBENTS Microsoft Copilot Studio enforces workflows but lacks flexibility

TIMING Enterprises starting to require auditable and rollbackable agents

4. AI-generated internal tools are low quality and dangerous • Heat 6/10 • AI Development

Writing an ops dashboard with AI is easy, but it might directly operate production databases

STRATEGIC Huge internal tools market, security moat determines if enterprises dare to use AI-generated internal tools

INCUMBENTS Retool, Airplane, Budibase don't do AI code filtering

TIMING Explosion of AI-generated internal tools, security incidents starting to appear

5. Smart home switch status falsely shows success • Heat 5/10 • AI Hardware

Mi Home shows light off, but it's still on

STRATEGIC Core smart home experience affects brand trust

INCUMBENTS Mi Home and Huawei Smart Life lack real status sync

TIMING Smart home adoption rising, false success complaints increasing

6. Apple Watch battery can't last long hike • Heat 4/10 • AI Hardware

Sports mode drops below 20% in 6 hours, needs mid-trip charge

STRATEGIC Core wearable pain point, battery directly impacts market share

INCUMBENTS Apple Watch and Galaxy Watch can't last full day heavy exercise

TIMING Outdoor sports population growing, competitor Garmin has battery advantage

// RADAR

Opportunity Radar (Solo & Small Team)

A 2×2 grid of demand strength vs. competition density.

Build now

Strong demand, low competition

- Colleagues waste time with AI-answerable questions
- Auto-generate step-by-step docs from screen recordings
- Notes scattered across four apps, impossible to find
- Distracted by editor while trying to focus on writing
- Zero-based men's skincare beginners can't find first step
- Friend borrows money without proof, hard to collect later

Watch closely

Strong demand, high competition

- Can't see which skills and tokens agent uses
- Code search grep is too primitive for agents
- AI agent reloads skills wasting context
- Fragmented lock file support across package managers
- Cross-platform CLI lacks unified DDC command support
- Wi-Fi router reconnection triggers light strip

Wait & validate

Weak demand, low competition

- Year of video footage lost on hard drive
- Video training for volunteers but VPS can't handle
- Content scraped by AI, no traffic back
- AI generates hollow content, wasting creativity
- Insurance AI gives product names but no selection help
- AI teaches outdated Rust syntax

Skip for now

Weak demand, high competition

- Unsure which grad school to choose
- No immersive practice for new language

// NEW & TRENDING

What's New This Period

First appearances

- Zero-based men's skincare beginners can't find first step
- AI teaches outdated Rust syntax

Rising fastest

- Multiple AI agents can't collaborate
- No security boundary for sensitive agent actions
- No reliable fallback when AI agent errors
- Can't see which skills and tokens agent uses
- Code search grep is too primitive for agents
- Year of video footage lost on hard drive
- Colleagues waste time with AI-answerable questions
- AI-generated internal tools are low quality and dangerous
- AI agent reloads skills wasting context
- Unsure which grad school to choose
- Auto-generate step-by-step docs from screen recordings
- Notes scattered across four apps, impossible to find
- Insurance AI gives product names but no selection help
- Apple Watch battery can't last long hike
- Video training for volunteers but VPS can't handle
- Content scraped by AI, no traffic back
- Cross-platform CLI lacks unified DDC command support

Cooling

- Smart home switch status falsely shows success

// USER VOICES

Selected User Voices

Real user comments from global tech community discussions.

"Using Claude Code, Codex, and Aider together—they each do their own thing, overwriting each other in the same project. I edit one file and all three agents jump in. So annoying."

— AI Development • Multiple AI agents can't collaborate

"Agents can execute commands, send emails, modify databases—all relying on prompt constraints. If the prompt gets bypassed, company data is toast."

— AI Development • No security boundary for sensitive agent actions

"Agents hit a wrong branch and retry in a loop, wasting tons of tokens and time. In the end, I have to kill the process manually."

— AI Development • No reliable fallback when AI agent errors

"Skills reload every session, context is never enough, and I have no idea which skills are unused—it's a total black box."

— AI Development • Can't see which skills and tokens agent uses

"Agent grep reads entire code, token usage explodes, and it often pulls in unrelated files. Searching for a function name can bring up the whole project."

— AI Development • Code search grep is too primitive for agents

"Recorded thousands of hours of video, but finding a specific scene relies on memory. Tons of footage on the hard drive, but I can never find what I need."

— AI Content • Year of video footage lost on hard drive

"Someone in the team asks me directly before even trying the AI. I've become a human search engine."

— AI Tools • Colleagues waste time with AI-answerable questions

"Using AI to write an ops panel is great, but it might directly operate the production database—scares me to death."

— AI Development • AI-generated internal tools are low quality and dangerous

// ABOUT

About This Report

Copyright & Sharing

© 2026 FindAI Go. Free to share with attribution to FindAI Go and link to findaigo.cn. For commercial use, please contact us first.

Contact

Website: findaigo.cn

Email: business@findaigo.com

Subscribe weekly: findaigo.cn/ai-radar

Disclaimer: This report is AI-generated and intended for product research and decision support. Decisions should be made based on your own judgment. FindAI Go is not liable for outcomes from decisions based on this report.