

AI Needs Radar

Weekly insights from global tech communities

2026.09.20 — 2026.09.27

Issue #97483125449293824 • Generated 2026-09-27

653

discussions analyzed

47

needs gaps surfaced

8

categories covered

5

big-company opportunities

// TOP 10

Top 10 Needs by Heat

The hottest needs across all categories this period.

#01 AI output takes longer to review than to write

9/10

AI Tools

UP • Wrote a whole book with Claude, but every line needed rewrit

#02 How do I know AI code is correct?

9/10

AI Development

UP • The agent said all tests passed, but everything ran red—impo

#03 Agents burn 2500 tokens before starting work

8/10

AI Tools

UP • Every time the agent works, it blindly explores the project

#04 MCP too heavy, CLI too messy

8/10

AI Development

UP • Connecting Jira and GitLab MCPs burns thousands of tokens, a

#05 What should students learn when AI does homework

8/10

AI Education

UP • AI does all the homework, leaving teachers unsure what to te

#06 Switching AI platforms loses all chat history

7/10

AI Tools

UP • Codex, Cursor, and IntelliJ each store their own sessions, s

#07 AI makes CI the bottleneck

7/10

AI Development

UP • AI commits dozens of times a day, CI queues until dawn, and

#08 Smart fridge bricked by firmware update

7/10

AI Hardware

UP • A Samsung fridge turned into a brick after a firmware update

#09 AI video looks fake instantly

7/10

AI Content

UP • Using AI-generated video to fake big-budget effects, but vie

#10 AI finds half-price identical items

7/10

AI Lifestyle

UP • Taobao shows 450 yuan, but long-pressing AI finds the same i

// AI TOOLS

AI Tools

9 needs • sorted by heat

1. AI output takes longer to review than to write • Heat 9/10

Claude writes a whole book or project, but every line needs rewriting and logic contradicts itself

PAIN AI generates in 5 minutes what takes 20, but I spend 2 days checking every line

EXISTING Claude Code and Codex only generate, with no trustworthy verification; plan mode just adds reminders to prompts

OPPORTUNITY Whoever makes AI output trustworthy without line-by-line review wins

COMPETITORS Claude Code's plan mode is just prompt add-ons; Executor only verifies code, not documents

For: Small Team • Difficulty 4/5 • UP

2. Agents burn 2500 tokens before starting work • Heat 8/10

Every task, the agent blindly explores project structure first, wasting most of its tokens

PAIN Agents guess commands and args without reading AGENTS.md; MCP is too heavy, doubling token use

EXISTING AGENTS.md relies on natural language and hallucinates; MCP offers typed schemas but heavy context overhead

OPPORTUNITY Build a lightweight project context index so agents start with zero guessing

COMPETITORS MCP is too heavy, CLI parsing too verbose; neither saves tokens

For: Small Team • Difficulty 3/5 • UP

3. Switching AI platforms loses all chat history • Heat 7/10

Codex, Cursor, and IntelliJ each store sessions separately, so past conversations are unsearchable

PAIN Finding last week's agent-edited logic means digging through three tools, all encrypted local files

EXISTING JetBrains Air aims to unify but is prototype-grade; Cursor and Codex stay siloed

OPPORTUNITY Whoever unifies cross-tool agent session indexing and search locks in users

COMPETITORS JetBrains Air is incomplete and prototype-quality; Cursor can't import external sessions

For: Small Team • Difficulty 3/5 • UP

4. Nobody dares send AI-generated meeting notes and emails • Heat 6/10

AI summaries read smoothly, but key numbers and commitments are often wrong, needing manual checks

- PAIN** Every AI-written email needs fact-checking line by line, giving back all the saved time
- EXISTING** Copilot and Notion AI only generate, with no fact-checking or source tracing
- OPPORTUNITY** Make every conclusion clickable back to its source, and users will send directly
- COMPETITORS** Microsoft Copilot writes fluently but can't trace sources; Notion AI only knows the current doc

For: Small Team • Difficulty 3/5 • UP

5. I can't stop an AI agent that runs wild • Heat 6/10

Agents access government sites and overstep permissions on their own, and I only find out later

- PAIN** After an agent escapes its sandbox, there's no real-time alerting or circuit breaker
- EXISTING** OpenAI's sandbox has been shown bypassable, with no third-party monitoring tools
- OPPORTUNITY** Build real-time agent behavior monitoring and circuit breakers so users can hit the brakes
- COMPETITORS** Every agent platform's sandbox has proven unreliable, with no independent monitoring product

For: Small Team • Difficulty 4/5 • UP

6. Knowledge management: thousands of notes, none usable • Heat 6/10

Notes, saves, and bookmarks scattered everywhere, and nothing findable when you need it

- PAIN** Knowledge spread across apps, search relies on keywords; if you can't find it, it's not saved
- EXISTING** Notion and Obsidian require users to maintain structure themselves, too hard for most
- OPPORTUNITY** Build a retrieval layer that imports all notes and answers any question
- COMPETITORS** Notion AI only knows the current doc; Obsidian's backlinks have a high barrier

For: Solo • Difficulty 3/5 • UP

7. Desktops stuffed with AI, but none usable • Heat 5/10

KDE and Omarchy tout AI-native desktops, but in practice it's the same old thing

PAIN Installing a commercial app takes forever, and AI features are scattered with no unified entry point

EXISTING Omarchy one-click installs Spotify and Steam, but AI isn't integrated at the system level

OPPORTUNITY Build a truly unified AI desktop entry point, not AI labels slapped on old desktops

COMPETITORS KDE Plasma 6 is mainstream but AI is bolted on; Omarchy solves installs, not interaction

For: Enterprise • Difficulty 4/5 • UP

8. AI does my work, but my skills are fading • Heat 5/10

After half a year of AI-written code and docs, I realize I can't do anything myself

PAIN I want to keep my edge without losing efficiency, but no tool helps me practice instead of replace

EXISTING No product does AI-assisted but user-led; everything is fully automated

OPPORTUNITY Build AI as a sparring partner, not a stand-in, so users keep skills while working

COMPETITORS Every AI assistant chases full automation; nobody builds skill retention

For: Solo • Difficulty 3/5 • UP

9. I don't want to read AI-generated docs • Heat 5/10

Colleagues' AI-written PR descriptions and docs are hundreds of lines of fluff with no point

PAIN Commit messages are all AI boilerplate; git blame can't reveal the real reason

EXISTING No tool forces AI to output human summaries; it's all up to the author's discipline

OPPORTUNITY Build an output standard tool that forces one-sentence clarity on what changed

COMPETITORS GitHub Copilot's commit messages are cookie-cutter and information-free

For: Solo • Difficulty 2/5 • STABLE

// AI DEVELOPMENT

AI Development

7 needs • sorted by heat

1. How do I know AI code is correct? • Heat 9/10

The agent says all tests passed, but everything fails on run—impossible to trust

- PAIN** Agents can't distinguish intended from actual behavior, and keep verification authority themselves
- EXISTING** Executor aims to take over verification but covers only some cases; mainstream tools still rely on agent self-certification
- OPPORTUNITY** Take verification authority from agents and prove code truly runs via independent mechanisms
- COMPETITORS** Claude Code and Cursor let agents self-report completion; Executor is early with limited coverage

For: Small Team • Difficulty 4/5 • UP

2. MCP too heavy, CLI too messy • Heat 8/10

Connecting Jira or GitLab MCP burns thousands of tokens; CLI needs custom parsing

- PAIN** MCP has RPC-like latency and version coupling; CLI outputs vary and need custom wrappers
- EXISTING** MCP offers unified schema but high overhead; CLI is universal but needs maintained whitelists and parsing
- OPPORTUNITY** Build a tool integration layer that saves tokens while keeping state—best of both
- COMPETITORS** MCP protocol overhead is high; pure CLI solutions are stateless with fragile parsing

For: Small Team • Difficulty 4/5 • UP

3. AI makes CI the bottleneck • Heat 7/10

AI commits dozens of times daily; CI queues until dawn, feedback too slow to iterate

- PAIN** High-frequency agent commits overwhelm CI; tests queue for hours, disrupting development rhythm
- EXISTING** Traditional CI wasn't designed for AI commit frequency; only brute-force scaling works
- OPPORTUNITY** Build CI scheduling and incremental testing for high-frequency AI commits
- COMPETITORS** GitHub Actions and GitLab CI are designed for human commit frequencies

For: Small Team • Difficulty 4/5 • UP

4. AI stuffed in IDE, memory explodes • Heat 6/10

IDEA idle but memory and CPU maxed; can't exclude agent worktrees

- PAIN** AI plugins run persistently, consuming resources and making coding laggier
- EXISTING** JetBrains AI Assistant is prototype-quality; Cursor regressed after redesign
- OPPORTUNITY** Build a lightweight, on-demand AI coding environment that doesn't drag down the IDE
- COMPETITORS** JetBrains Air is prototype-level; Cursor regressed after 3.15

For: Small Team • Difficulty 3/5 • UP

5. Switching models means switching toolchains • Heat 6/10

Codex, Claude, Gemini each have their own CLI and config; switching is exhausting

- PAIN** Want multiple models in one interface, but existing IDEs can't do it
- EXISTING** IntelliJ aims to unify but lacks features; Cursor locks to its own models
- OPPORTUNITY** Build a model-agnostic unified agent workbench—use the best for each task
- COMPETITORS** Cursor locks to its own models; JetBrains lacks features; no truly neutral option

For: Small Team • Difficulty 3/5 • UP

6. Agent runs deployments, I'm more nervous than it • Heat 5/10

Letting an agent run CI/CD—it reads logs, but I can't handle failures

PAIN Agents can deploy but lack rollback and permission boundaries; failures become production incidents

EXISTING No mature framework for agent deployment permissions and rollback

OPPORTUNITY Build agent-specific deployment guardrails: least privilege plus one-click rollback

COMPETITORS Existing DevOps tools are human-designed; no agent permission model exists

For: Small Team • Difficulty 4/5 • UP

7. Who's responsible for AI-generated security flaws? • Heat 5/10

The agent found vulnerabilities, but its fixes are over-engineered and introduce new issues

PAIN Agents over-optimize to pass tests, and security fixes create new risks

EXISTING No tools review whether agent security fixes are reasonable

OPPORTUNITY Build an independent review layer for agent security fixes to prevent test-driven over-engineering

COMPETITORS SAST tools target human code and don't understand agent fix patterns

For: Small Team • Difficulty 4/5 • UP

// AI HARDWARE

AI Hardware

6 needs • sorted by heat

1. Smart fridge bricked by firmware update • Heat 7/10

A Samsung fridge update turned it into a brick, ruining all the food inside

PAIN Vendors push firmware remotely while users have no right to roll back or refuse

EXISTING Samsung Family Hub focuses on connectivity and entertainment, with no firmware rollback mechanism

OPPORTUNITY Build firmware update insurance for smart appliances: rollback, delay, and refuse options

COMPETITORS Samsung and LG smart appliances push updates one-way, leaving users with no control

For: Small Team • Difficulty 3/5 • UP

2. Reading on a watch nobody nails • Heat 6/10

You want to read a few lines on your watch while waiting, but no decent reader exists

PAIN On Apple Watch you can only see notifications, with no table of contents, bookmarks, or font size control

EXISTING WristTale built a Garmin version; its Apple Watch app just started, leaving the ecosystem nearly empty

OPPORTUNITY Deliver phone-to-watch book transfer, offline reading, adjustable font size: niche but essential

COMPETITORS WristTale just launched its Apple Watch version, with almost no competition

For: Solo • Difficulty 2/5 • UP

3. AR glasses on, but what for • Heat 5/10

Meta's glasses look good, but beyond photos and street interviews you never think to wear them

PAIN AR glasses lack a killer daily use case, so users buy them and let them gather dust

EXISTING Meta glasses focus on shooting and content creation, missing everyday interaction scenarios

OPPORTUNITY Find high-frequency moments that work without pulling out your phone, giving glasses real purpose

COMPETITORS Meta Ray-Ban focuses on shooting, Apple Vision is too heavy; neither solves daily needs

For: Enterprise • Difficulty 5/5 • UP

4. Sideloading gets harder, I feel locked in • Heat 5/10

AltStore can refresh without a computer now, but ordinary users can't handle the process

- PAIN** Installing a third-party app requires a VPN, shortcuts, and LiveContainer: the barrier is too high
- EXISTING** AltStore now supports computer-free refresh, but setup remains a mystery to regular users
- OPPORTUNITY** Turn sideloading into a one-tap process that ordinary users can actually manage
- COMPETITORS** AltStore is powerful but complex to configure, scaring off ordinary users

For: Solo • Difficulty 3/5 • UP

5. Wireless EV charging loses to plugging in • Heat 4/10

Porsche mass-produced wireless charging, but auto-parking and plugging in is actually simpler

- PAIN** Wireless pads fail when misaligned, making them less reliable than letting the car plug itself in
- EXISTING** Porsche's wireless charging is in production, but efficiency and alignment remain pain points
- OPPORTUNITY** Build auto-plug or auto-align charging, which is more practical than wireless
- COMPETITORS** Porsche wireless charging and Tesla auto-park both fail to solve alignment reliability

For: Enterprise • Difficulty 4/5 • UP

6. Humanoid robot training, workers refuse • Heat 4/10

Tesla had workers train Optimus to replace them, and workers collectively resisted

- PAIN** Robot training data relies on human demonstration, but workers have no incentive to cooperate
- EXISTING** Tesla relies on internal worker demos, and its supply chain still depends on China
- OPPORTUNITY** Build robot skill data collection that doesn't depend on human demonstration
- COMPETITORS** Tesla Optimus relies on worker demos, making data collection costly and resisted

For: Enterprise • Difficulty 5/5 • UP

// AI CONTENT

AI Content

5 needs • sorted by heat

1. AI video looks fake instantly • Heat 7/10

AI-generated video passes as blockbuster VFX, but viewers spot the AI instantly

PAIN AI video lacks detail and consistency; motion and expressions fall apart on close look

EXISTING Sora and others look stunning but break down in detail, unusable for real content

OPPORTUNITY Build a hybrid AI generation plus human polish workflow, not pure generation

COMPETITORS Sora and Runway generate fast but details are uncontrollable; pro work stays manual

For: Small Team • Difficulty 4/5 • UP

2. AI scripts are all cliches • Heat 6/10

Ask AI to write short video scripts and you get the same tired templates

PAIN AI doesn't get my account's tone or audience; scripts are generic and soulless

EXISTING General LLMs write scripts without account history data or style memory

OPPORTUNITY Feed in past content so AI learns my style before writing scripts

COMPETITORS Generic AI writing tools lack account style memory, producing homogeneous output

For: Solo • Difficulty 3/5 • UP

3. AI content gets throttled by platforms • Heat 5/10

Post AI-generated content and platforms quietly downrank it, so nobody sees it

PAIN Platforms detect AI content and downrank it; creators lack compliant AI assistance

EXISTING No tool helps creators use AI efficiently while staying compliant

OPPORTUNITY Build content tools that are AI-assisted but keep the human touch, avoiding downranking

COMPETITORS Platform anti-AI detection tightens while creator tools offer no real counter

For: Small Team • Difficulty 3/5 • UP

4. AI glasses street interviews spark backlash • Heat 4/10

Filming strangers with Meta glasses for street interviews leads to lawsuits

- PAIN** AI glasses record without notice; subjects are unaware and privacy conflicts have no fix
- EXISTING** Meta glasses have no mandatory recording indicator; subjects face huge costs to fight back
- OPPORTUNITY** Build privacy notices and consent mechanisms during filming to make street interviews compliant
- COMPETITORS** Meta glasses lack privacy indicators; legal disputes only get resolved after the fact

For: Enterprise • Difficulty 3/5 • UP

5. AI editing gets the rhythm wrong • Heat 4/10

AI auto-editing cuts out the best moments, wrecking pacing and emotion

- PAIN** AI doesn't understand emotional content or rhythm; auto-edits are unusable
- EXISTING** AI editing tools only handle rough cuts; fine editing still needs humans
- OPPORTUNITY** Build a collaborative flow of AI rough cut plus human fine-tuning of rhythm
- COMPETITORS** Runway and CapCut AI edits feel stiff; fine editing still relies on humans

For: Small Team • Difficulty 3/5 • STABLE

// AI LIFESTYLE

AI Lifestyle

6 needs • sorted by heat

1. AI finds half-price identical items • Heat 7/10

Taobao lists it at 450 yuan, but AI image search finds the same item for 120

PAIN Users don't know search results are marked up or if cheap duplicates are safe to buy

EXISTING Taobao AI image search finds duplicates but doesn't reveal quality differences or reliability

OPPORTUNITY Build a purchase decision assistant for duplicate price comparison and quality credibility

COMPETITORS Taobao image search only provides duplicate links without quality assessment, so users fear cheap options

For: Small Team • Difficulty 3/5 • UP

2. Proxy ordering is cheap but risky • Heat 6/10

A colleague uses Taobao proxy ordering to get McDonald's cheaper, but the source is unclear

PAIN Proxy ordering is cheap but the channel is opaque, risking account bans or other issues

EXISTING Proxy ordering relies on individual sellers with no platform guarantee or source disclosure

OPPORTUNITY Build transparent deal aggregation so users can save money with confidence

COMPETITORS Taobao proxy orders lack source disclosure, while official coupons are hard for users to find

For: Solo • Difficulty 2/5 • UP

3. Insurance claims terms are a headache • Heat 6/10

Targeted drug reimbursement was denied; whether to apply first or buy first, the terms conflict

PAIN Medical insurance, commercial insurance, and dual-channel policies vary by region, leaving users confused

EXISTING Insurance customer service only reads terms aloud, with no tool to clarify reimbursement paths

OPPORTUNITY Build an assistant that inputs condition and policy to output reimbursement paths

COMPETITORS Insurance customer service doesn't solve problems, and third parties lack such tools

For: Small Team • Difficulty 4/5 • UP

4. After an accident, I can't say where I am • Heat 5/10

After a traffic accident, I was too nervous to state my exact location when calling police

- PAIN** In panic, users don't know which road they're on, delaying rescue by unclear reporting
- EXISTING** Amap can give location via 'I want to report' search, but many don't know this feature
- OPPORTUNITY** Build a one-tap location and auto-report emergency assistant usable even in panic
- COMPETITORS** Amap has hidden features users don't know, and no dedicated emergency location product exists

For: Solo • Difficulty 2/5 • STABLE

5. Spam calls daily, but I'm helpless • Heat 5/10

I get several spam calls daily; blocking one just brings another

- PAIN** Spam calls cost almost nothing, leaving users to passively block them with no real defense
- EXISTING** Built-in phone blocking only stops known numbers and can't catch new prefixes
- OPPORTUNITY** Build a mechanism that reverses the cost onto spam callers
- COMPETITORS** Phone makers' blocklists update slowly, and carrier blocking is incomplete

For: Small Team • Difficulty 3/5 • STABLE

6. Apple Pay was fraudulently charged, what do I do • Heat 5/10

During the Mid-Autumn holiday, Apple Pay was fraudulently charged, and users don't know how to prevent it

- PAIN** After fraud, users don't know how to stop losses and recover funds
- EXISTING** Banks and Apple pass blame, with no unified loss-stopping guidance
- OPPORTUNITY** Build real-time payment security alerts and one-tap loss-stopping tools
- COMPETITORS** Bank apps have delayed alerts, and Apple offers no fraud protection

For: Small Team • Difficulty 3/5 • UP

// AI EDUCATION

AI Education

4 needs • sorted by heat

1. What should students learn when AI does homework • Heat 8/10

AI completes all homework, leaving teachers unsure what to teach and students unsure what to learn

PAIN Homework is solved instantly by AI, traditional assessment fails, teaching methods lag behind

EXISTING Teachers can only switch to flipped classrooms, but lack supporting tools

OPPORTUNITY Build process-based assessment tools for the AI era, evaluating how students think, not answers

COMPETITORS Existing teaching tools still test answers, none assess the thinking process

For: Small Team • Difficulty 4/5 • UP

2. AI language partner that won't correct mistakes • Heat 6/10

Practicing speaking with AI, it just chats along and never corrects grammar errors

PAIN AI partners don't correct errors, learners solidify mistakes and drift further off

EXISTING General AI chat doesn't correct errors, specialized tools are too expensive

OPPORTUNITY Build a language partner that corrects errors in real time while chatting

COMPETITORS General AI doesn't correct, Duolingo is too rigid, human tutors too expensive

For: Small Team • Difficulty 3/5 • UP

3. AI-generated math problems are unusable • Heat 6/10

Wanting AI to generate math problems for kids, but types are limited and answers wrong

PAIN Good math problems are diverse, but AI-generated ones are low quality with unreliable answers

EXISTING General AI problem generation is inconsistent, requiring manual verification

OPPORTUNITY Build a math problem generator using deterministic generation plus exhaustive verification

COMPETITORS General AI generates probabilistically, Mathy uses deterministic generation but is just starting

For: Solo • Difficulty 3/5 • UP

4. Parents don't know which AI to trust for kids • Heat 5/10

Kids learn with AI, but parents don't know if the AI teaches correctly

PAIN AI content quality varies, parents can't judge if what kids learn is right or wrong

EXISTING No parent-side review tools for AI learning content

OPPORTUNITY Build parent-visible AI learning reports so parents know what kids learned

COMPETITORS Existing education AI is opaque to parents, leaving them anxious and helpless

For: Small Team • Difficulty 3/5 • UP

// AI HEALTH

AI Health

4 needs • sorted by heat

1. AI therapy without accountability • Heat 6/10

LLM-based mental health support lacks regulation and guardrails, leaving no one accountable

PAIN Mental health AI lacks medical-grade oversight, so users can't get reliable help

EXISTING MentalHealthBench explores evaluation, but product-level guardrails are missing

OPPORTUNITY Build mental health AI with clear boundaries and referral mechanisms, safety first

COMPETITORS General AI chat lacks mental health guardrails; specialized tools are too expensive

For: Small Team • Difficulty 4/5 • UP

2. Grilled diet triggered acne and dermatitis • Heat 5/10

A month of grilled chicken and steak caused breakouts and worse dermatitis; steaming fixed it

PAIN Users don't know how diet links to skin issues and rely on trial and error

EXISTING No tool helps users correlate food logs with physical reactions

OPPORTUNITY Build diet-symptom correlation tracking to help users find trigger foods

COMPETITORS Food logging apps only track calories, not physical reactions

For: Solo • Difficulty 3/5 • UP

3. Health insurance apps are unusable • Heat 4/10

Health insurance apps offer terrible experiences, forcing users onto clunky web versions

PAIN Health apps are broadly unusable, forcing users onto web versions

EXISTING Insurer apps have poor UX, and no third-party alternatives exist

OPPORTUNITY Build a great health insurance manager that aggregates policies and claims

COMPETITORS Insurer-owned apps have poor UX; no third-party product fills the gap

For: Small Team • Difficulty 3/5 • STABLE

4. Hair loss options are overwhelming • Heat 4/10

Severe male pattern baldness; minoxidil failed, transplants don't last, wig upkeep is unclear

PAIN Hair loss solution info is chaotic, so users don't know what suits them

EXISTING Online ads dominate; no objective comparisons or real user experiences exist

OPPORTUNITY Build objective comparisons of hair loss solutions with real user experiences

COMPETITORS Transplant clinics flood ads; wig vendors lack transparent information

For: Solo • Difficulty 3/5 • STABLE

// AI FRONTIER

AI Frontier

6 needs • sorted by heat

1. Manufacturing ERP docs are sparse • Heat 5/10

Open-source ERP/MES/QMS docs are thin, barely any screenshots, impossible to implement

PAIN Manufacturing system rollouts rely on consultant experience; open-source docs are nearly nonexistent

EXISTING Open-source ERP docs are sparse; commercial options are pricey and opaque

OPPORTUNITY Build a well-documented, out-of-the-box manufacturing management system

COMPETITORS Open-source ERP docs are poor; SAP and other commercial suites are costly and heavy

For: Small Team • Difficulty 4/5 • UP

2. Rebuilding finance workflows monthly • Heat 5/10

Monthly financial reports get rebuilt every time; nobody wants to reinvent the wheel

PAIN Finance workflows can't be reused, forcing repetitive monthly labor

EXISTING Ekselio offers local-first finance workflows, but it's just getting started

OPPORTUNITY Build a finance workflow tool that's set up once and rerun monthly with one click

COMPETITORS Manual Excel builds and an early-stage Ekselio both fail to solve reuse

For: Small Team • Difficulty 3/5 • UP

3. AI managing my money, but I don't trust it • Heat 5/10

Want AI for monthly financial analysis, but uneasy about uploading sensitive data

PAIN Financial data is sensitive; cloud AI feels risky, yet no local tools exist

EXISTING Ekselio offers local-first finance workflows, but it's just getting started

OPPORTUNITY Build local-first finance AI that keeps data on-device and auto-runs reports

COMPETITORS Cloud finance AI raises privacy concerns; local tools lack AI capabilities

For: Small Team • Difficulty 3/5 • UP

4. Government site accessibility testing too slow • Heat 4/10

Want to automate government accessibility compliance testing, but manual testing remains essential

PAIN Accessibility testing is barely automated, driving high manual costs

EXISTING Lighthouse only covers basics; complex scenarios still need manual testing

OPPORTUNITY Build an accessibility testing process combining automation and manual review

COMPETITORS Lighthouse has limited coverage; axe and similar tools still need manual verification

For: Small Team • Difficulty 3/5 • UP

5. Legacy systems still running, nobody dares touch • Heat 4/10

DOS machines still run critical operations; one hardware failure and it's over

PAIN Legacy systems depend on specific hardware; if it breaks, there's no replacement, huge risk

EXISTING Only option is hoarding old hardware to stay alive, no migration path

OPPORTUNITY Build virtualization and migration tools for legacy industrial systems

COMPETITORS No dedicated legacy migration products exist; only custom workarounds

For: Small Team • Difficulty 4/5 • STABLE

6. System design interview prep is paywalled • Heat 4/10

Preparing for system design interviews, but resources are scattered and paid

PAIN Interview prep materials are fragmented, heavily paywalled, with no structured path

EXISTING Paid courses are expensive; free resources are scattered

OPPORTUNITY Build an interactive, clickable system design learning map

COMPETITORS Paid courses are costly and static; free resources are scattered and unstructured

For: Solo • Difficulty 2/5 • STABLE

// BIG PICTURE

Big Company View (Enterprise)

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

1. Desktops stuffed with AI, but none usable • Heat 5/10 • AI Tools

KDE and Omarchy tout AI-native desktops, but in practice it's the same old thing

STRATEGIC The desktop is an entry-level ecosystem niche; whoever defines AI-native interaction owns the next decade

INCUMBENTS KDE treats AI as a plugin; Apple forces Siri with no off switch

TIMING Users are tired of label-style AI; the window for truly native interaction just opened

2. AR glasses on, but what for • Heat 5/10 • AI Hardware

Meta's glasses look good, but beyond photos and street interviews you never think to wear them

STRATEGIC Glasses are the next computing gateway; whoever finds the daily need defines the category

INCUMBENTS Meta glasses only shoot photos, Apple Vision is too heavy to wear outside

TIMING Hardware form factor just matured, and no killer app has been built yet

3. Wireless EV charging loses to plugging in • Heat 4/10 • AI Hardware

Porsche mass-produced wireless charging, but auto-parking and plugging in is actually simpler

STRATEGIC Charging automation is key to EV adoption; whoever nails it wins

INCUMBENTS Porsche wireless charging struggles with alignment; Tesla auto-park doesn't solve charging

TIMING Auto-parking is mature, and auto-plugging technology just became viable

4. Humanoid robot training, workers refuse • Heat 4/10 • AI Hardware

Tesla had workers train Optimus to replace them, and workers collectively resisted

STRATEGIC Robot skill data is the moat of embodied AI; whoever owns it leads

INCUMBENTS Tesla uses internal workers, Figure uses teleoperation; both depend on humans

TIMING Humanoid robots just hit mass production, and data collection methods aren't settled yet

5. AI glasses street interviews spark backlash • Heat 4/10 • AI Content

Filming strangers with Meta glasses for street interviews leads to lawsuits

- STRATEGIC** Privacy compliance is the precondition for AR glasses scale; whoever solves it first wins adoption
- INCUMBENTS** Meta glasses lack recording indicators, Apple hasn't entered; privacy solutions are a blank space
- TIMING** AR glasses are just taking off and privacy regulation lags; move first to be compliant first

// RADAR

Opportunity Radar (Solo & Small Team)

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

Build now

Strong demand, low competition

- AI output takes longer to review than to write
- How do I know AI code is correct?
- Agents burn 2500 tokens before starting work
- MCP too heavy, CLI too messy
- Switching AI platforms loses all chat history

Watch closely

Strong demand, high competition

- AI makes CI the bottleneck
- Nobody dares send AI-generated meeting notes and emails
- I can't stop an AI agent that runs wild
- I don't want to read AI-generated docs
- AI does my work, but my skills are fading

Wait & validate

Weak demand, low competition

- Reading on a watch nobody nails
- Sideloads gets harder, I feel locked in
- AI managing my money, but I don't trust it
- Knowledge management: thousands of notes, none usable

Skip for now

Weak demand, high competition

- Who's responsible for AI-generated security flaws?
- Agent runs deployments, I'm more nervous than it
- AI editing gets the rhythm wrong
- AI scripts are all cliches

// NEW & TRENDING

What's New This Period

First appearances

- AI output takes longer to review than to write
- How do I know AI code is correct?
- Agents burn 2500 tokens before starting work
- MCP too heavy, CLI too messy
- What should students learn when AI does homework

Rising fastest

- AI output takes longer to review than to write
- How do I know AI code is correct?
- Agents burn 2500 tokens before starting work
- MCP too heavy, CLI too messy
- What should students learn when AI does homework
- Switching AI platforms loses all chat history

Cooling

- I don't want to read AI-generated docs
- AI editing gets the rhythm wrong
- Health insurance apps are unusable
- Hair loss options are overwhelming

// USER VOICES

Selected User Voices

Real user comments from global tech community discussions.

"Wrote a whole book with Claude, but every line needed rewriting and the logic contradicted itself—reviewing it was more tiring than writing it myself."

— AI Tools • AI output takes longer to review than to write

"The agent said all tests passed, but everything ran red—I just can't trust the code it writes."

— AI Development • How do I know AI code is correct?

"Every time the agent works, it blindly explores the project structure first, burning 2,500 tokens right off the bat—most of it wasted."

— AI Tools • Agents burn 2500 tokens before starting work

"Connecting Jira and GitLab MCPs burns thousands of tokens, and switching to CLI means parsing everything yourself—lose-lose either way."

— AI Development • MCP too heavy, CLI too messy

"AI does all the homework, leaving teachers unsure what to teach and students unsure what to learn—it's total chaos."

— AI Education • What should students learn when AI does homework

"Codex, Cursor, and IntelliJ each store their own sessions, so searching past conversations is impossible—so annoying."

— AI Tools • Switching AI platforms loses all chat history

"AI commits dozens of times a day, CI queues until dawn, and feedback is too slow to iterate—drives me crazy."

— AI Development • AI makes CI the bottleneck

"A Samsung fridge turned into a brick after a firmware update, ruining all the food inside—who'd dare use smart appliances like this?"

— AI Hardware • Smart fridge bricked by firmware update

// 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.