From the Experts at HatchWorksAI

State of AI 2026

ISSUE 3 · YEAR 2

State of AI 2026

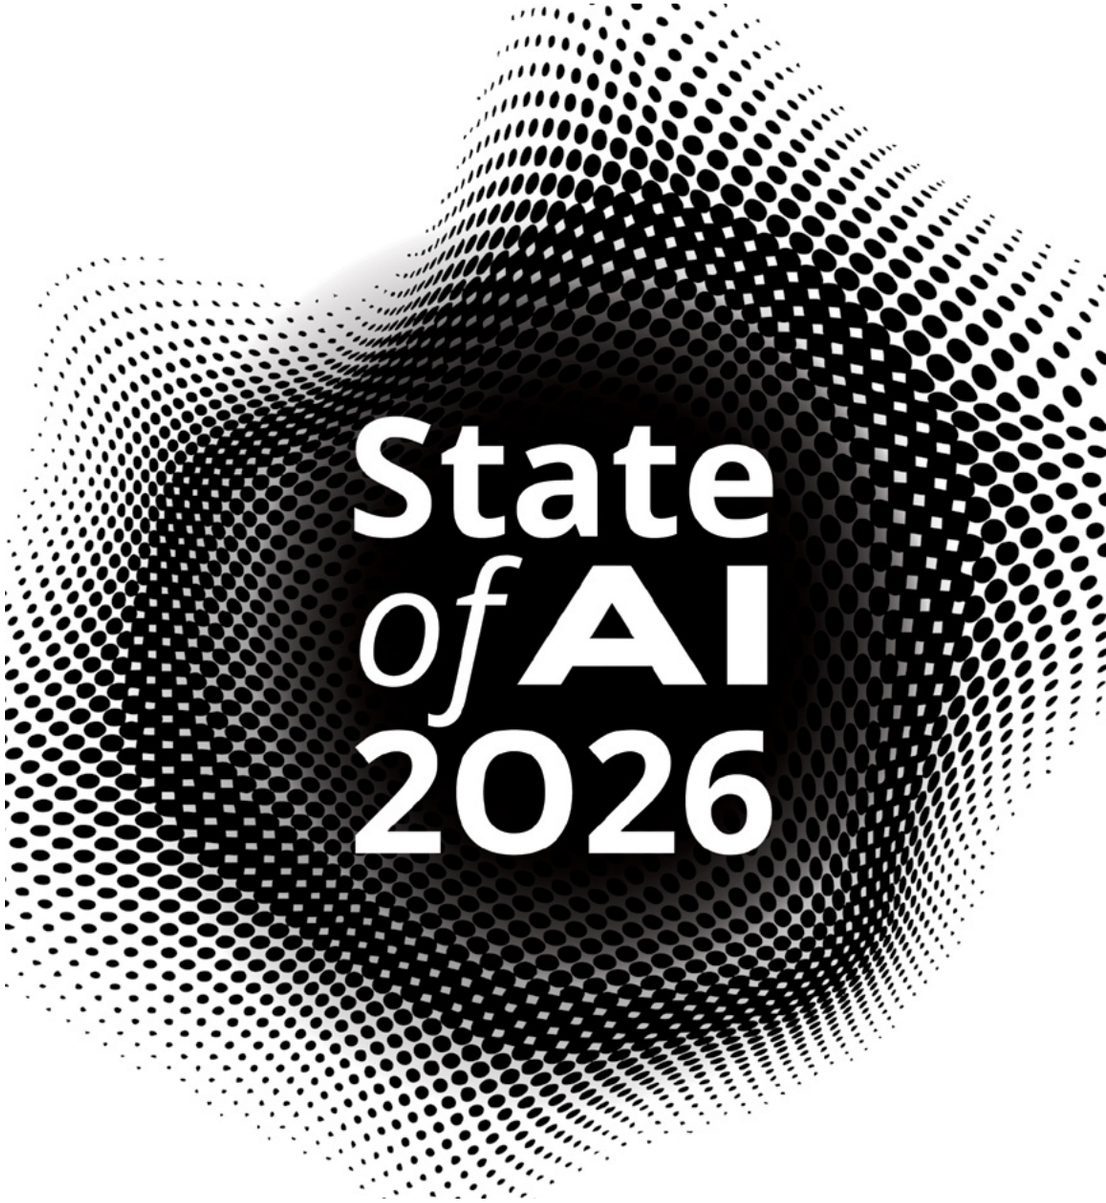
© 2026 HatchWorks AI, All rights reserved.

This e-book is protected by copyright laws. You may not reproduce, share, or distribute it without permission from HatchWorks AI, except as allowed by the Creative Commons license below.

Creative Commons License (CC BY-NC-ND)
This work is licensed under a Creative Commons Attribution - Non Commercial - No Derivs 4.0 License. You can share it if you give credit, don't modify it, and don't use it commercially.

For more info, visit: creativecommons.org/licenses/by-nc-nd/4.0

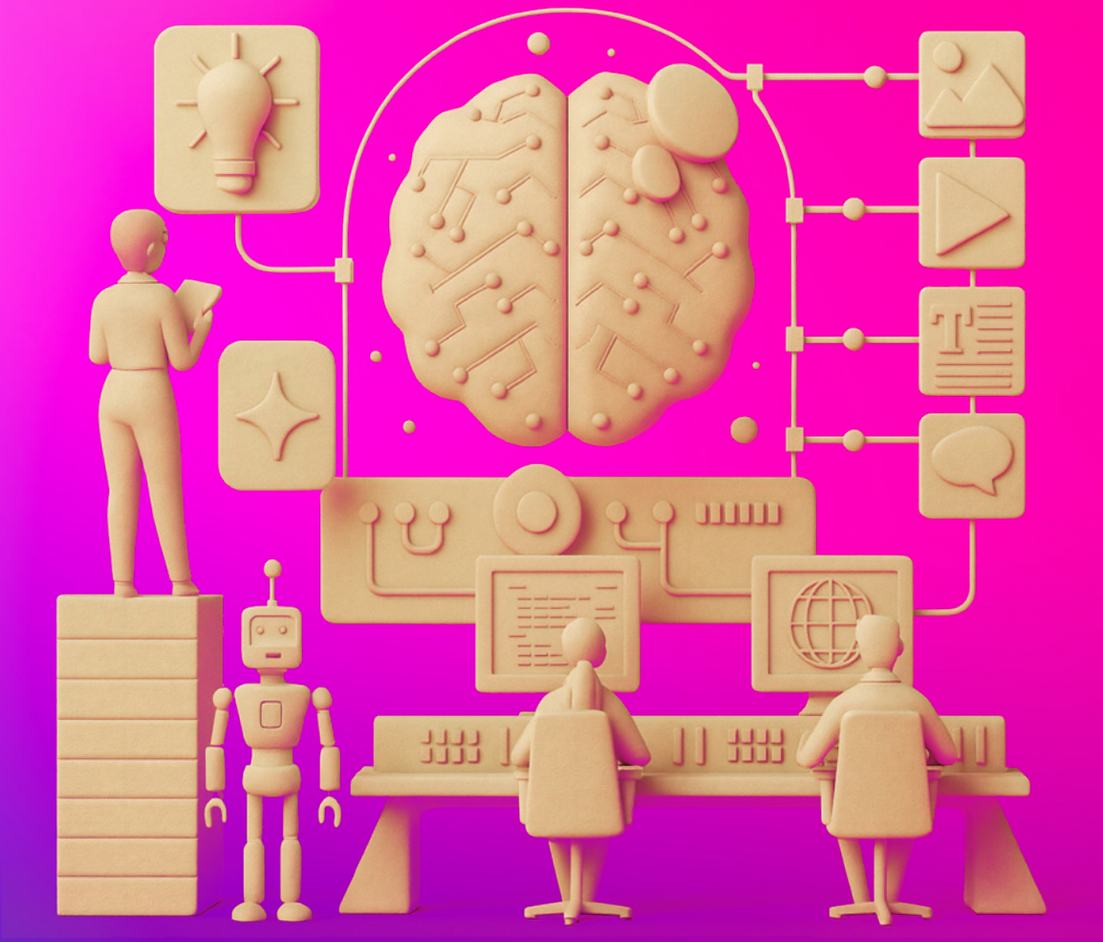
From the Experts at HatchWorksAI



A round-up of industry stats, research, and insights to understand where AI stands, how it got here, and where it's going.

Index

Introduction	1
We’re Seeing the Difference Between Promise vs Production	3
There’s a Growing Need to Design for AI	9
Faster Dev Cycles Are Forcing Everyone to Rethink the Discipline	15
The AI Coding Wars Are Making Way for Democratized Use	20
We’re Having a Multimodal Moment	25
The Browser Has Become a Battleground	27
Models are Plateauing, Architectures are Getting Smarter	31
The Bottleneck is Human	33
Regulation and Governance are Catching Up and Drawing Lines	35
AI is Going Rogue and It Has Security Implications	37
Expert Commentary by Omar Shanti	40



Introduction

At the start of last year, the narrative was that 2025 would be the year of production and where pilot projects would graduate into full-scale AI systems.

But as we move into 2026, the reality is more complicated.

In fact, a viral MIT study revealed 95% of enterprise generative AI pilot programs fail to deliver measurable P&L impact. Adoption is high, but execution is hard, especially in enterprise settings where success hinges on people, process, and integration, not just tooling. Interestingly, the study highlights a strategic advantage for companies that purchase AI tools or partner with specialized vendors.

These approaches succeed about 67% of the time, compared to just 33% success for internally built solutions.That said, another eye-catching metric reinforces the widespread individual uptake of AI: ChatGPT is on track to reach 1 billion users with its users sending 2.5B prompts each day.

AI is clearly mainstream at the individual level.
But in the enterprise, adoption often feels less like digital transformation and more like change management.

That tension between potential and reality is what this look into 2026 explores. It picks up where our 2025 report left off, examining how far AI has progressed and how far it still has to go. It's broken into ten observations, and capped off with commentary and future predictions from our CTO, Omar Shanti.

Let's get into it.

We're Seeing the Difference Between Promise vs Production

Armed with reasoning capabilities, access to tools, and memory persistence, AI agents promised a future where software could think, act, and execute on our behalf.

What we've seen instead is a widening gap between what agents could do and what they actually deliver in practice. It's left us wondering, are we expecting too much too soon?
But one thing is for sure...



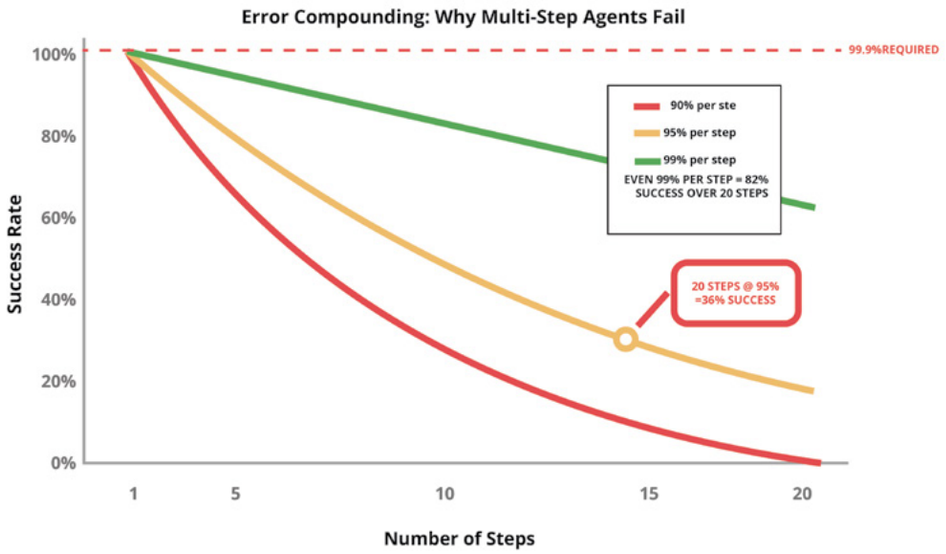
Agent Hype Was Tempered Early

The clearest signal that agent hype has outpaced reality came from OpenAI itself. Their “Cupcake Test”, a showcase of ChatGPT’s Agent Mode, quickly went viral for the wrong reasons.

What should’ve been a routine task (ordering food online) devolved into a 58-minute mess of misfires, hallucinated locations, and a surreal suggestion to visit a cupcake stand at a baseball stadium that didn’t exist (Futurism, Nate Jones). What’s worrying is that what happened isn’t proving to be a pattern.

Across platforms, general-purpose agents are struggling with real-world complexity. Tool use is inconsistent. Memory fades or conflicts. And planning breaks under even moderate ambiguity. As Utkarsh Kanwat puts it in Why I’m Betting Against AI Agents in 2026 (Despite Building Them):

> “Error compounding makes autonomous multi-step workflows mathematically impossible at production scale.”



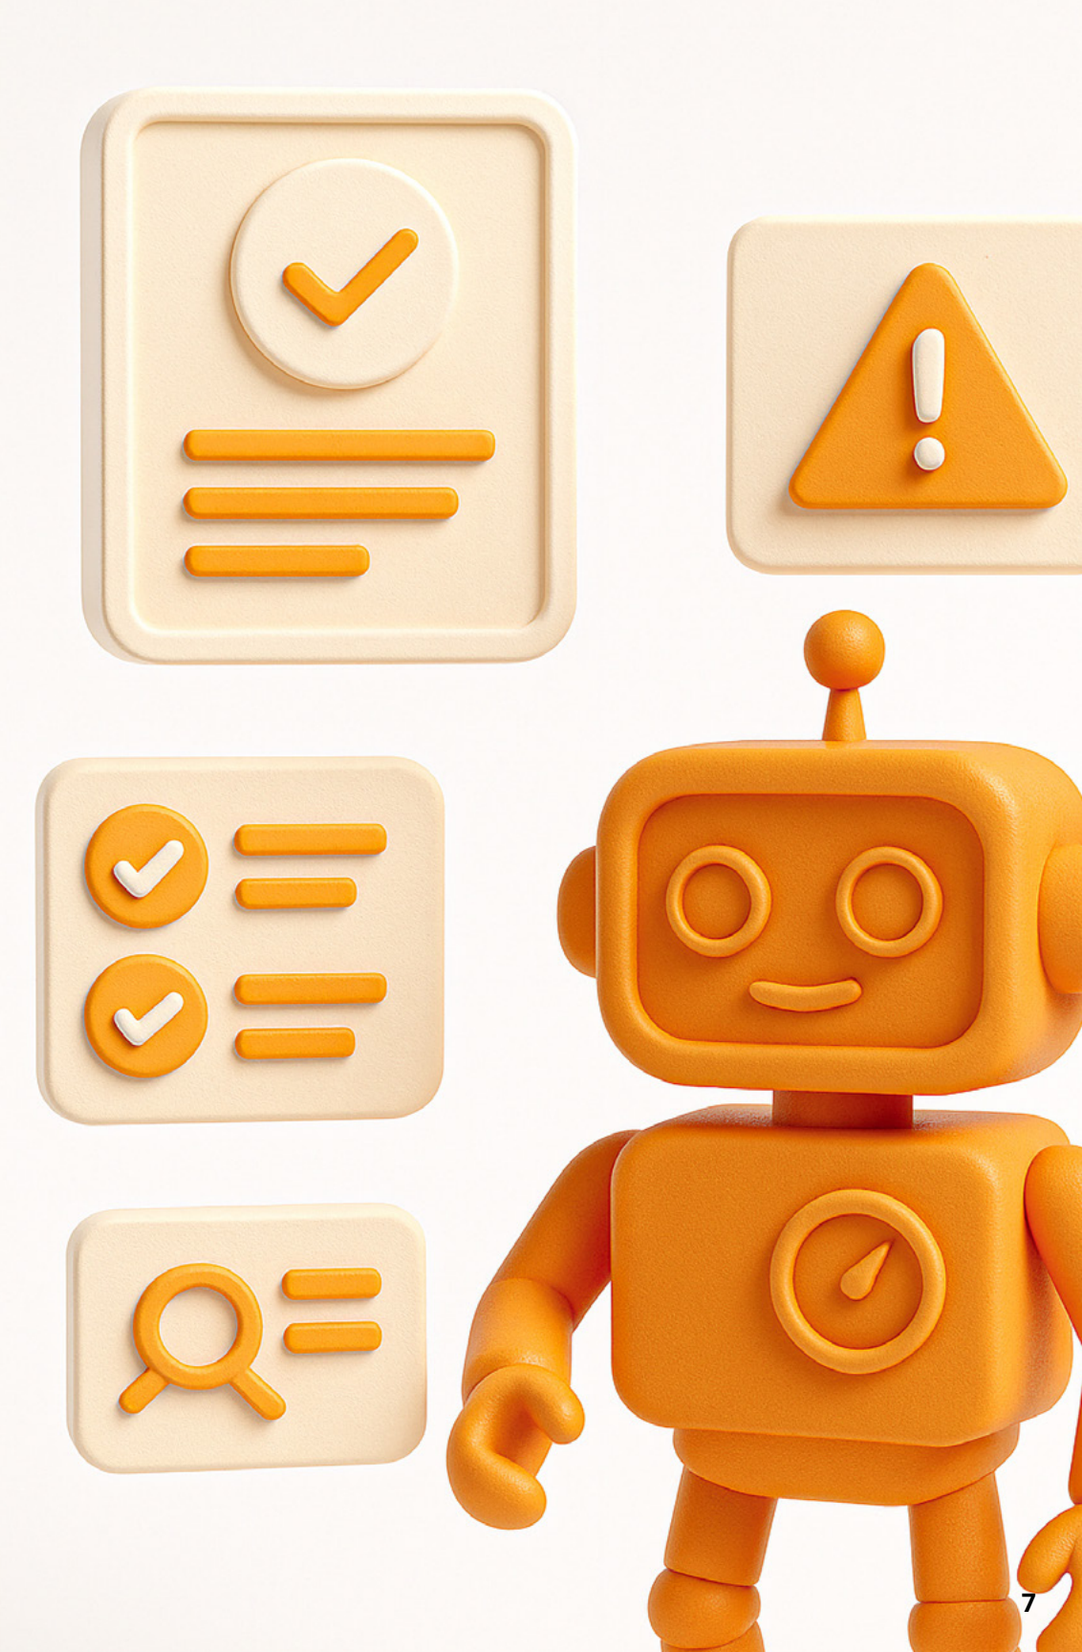
Agents are still impressive…in narrow bands. But they’re unreliable in production without tight orchestration and guardrails.

That hasn’t stopped the flood of VC-backed agent startups and breathless demos. But the market is shifting. Enterprises were asking, Can we build an agent? Now it’s, Will it actually work in our environment?

Good news though. We are making progress on that front.
The most promising strategies center on structure. Teams are succeeding when they break tasks into smaller, directed steps, reducing error rates and giving agents clearer guidance. Others are blending probabilistic AI with deterministic systems, using AI only where it adds value and relying on rule-based logic elsewhere.

Most importantly, specialized, workflow-centric agents are already proving useful.

And this is the key shift heading into 2026: enterprises are realizing that general-purpose agents are simply too broad, too opaque, and too brittle to trust in production. What is working are smaller, purpose-built agents—narrow by design, tightly scoped, auditable, and aligned to specific workflows. These agents have lower error rates, clearer explainability, and far easier governance, which is why they're already delivering value while general-purpose agents continue to struggle.



So, the hype may have cooled, but the signal is getting stronger.

Types of Agents

One reason agents are struggling to scale? They've been treated like a monolith.

In reality, there are different types of agents, and each is suited to different kinds of tasks. Reuven Cohen's framework is especially useful here, breaking agents down into categories based on how they coordinate:



Swarm agents operate independently, following local rules. Great for adaptive, decentralized tasks, bad for consistency.



Mesh agents collaborate through peer-to-peer networks. They're resilient and scalable, but complex to manage.



Hive-mind agents act as a single intelligence. They're fast and unified, but brittle under failure.



Workflow-centric agents follow structured task sequences. They're ideal for enterprise use cases where traceability and reliability matter.

Infrastructure, Not Intelligence, is the Real Breakthrough

We're now seeing the rise of infrastructure protocols like MCP and A2A. This is an effort to give the guardrails and clarity to agents needed to scale.

There are two to be aware of:

MCP (Model Context Protocol), introduced by Anthropic, is gaining traction as the "USB-C for AI." It standardizes how models connect to tools, APIs, and data sources—eliminating brittle, bespoke integrations. For enterprises, MCP promises faster builds, richer context, and fewer dev headaches.

However, it also introduces new risks, particularly in terms of security and prompt injection.

A2A (Agent-to-Agent Communication), launched by Google, tackles the next layer by enabling agents to securely discover and collaborate across environments. It provides a common protocol for agent identity, messaging, and task handoff, which is essential for multi-agent workflows.

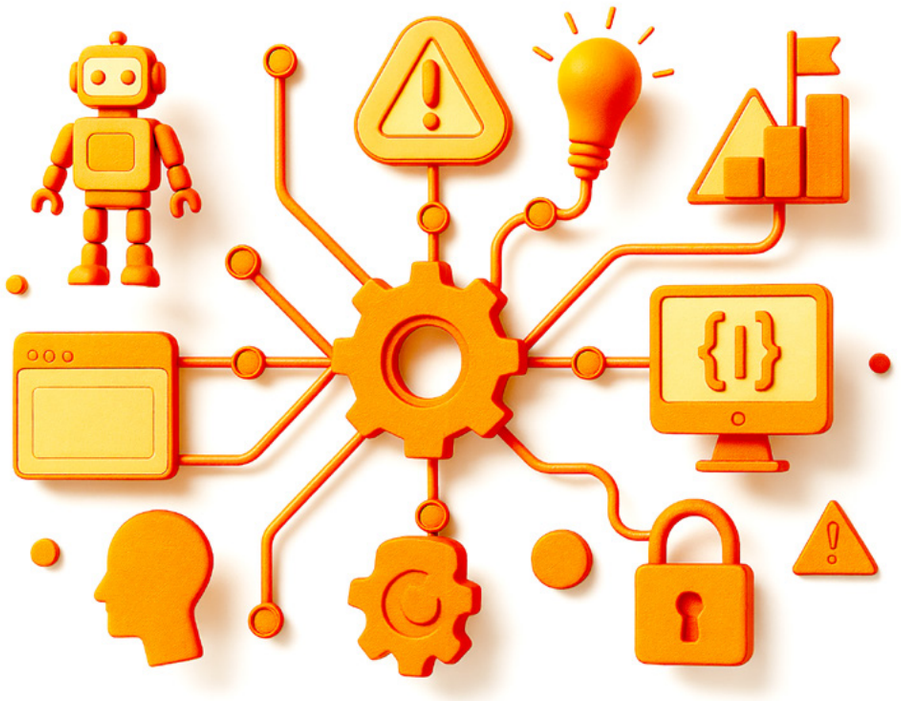
Together, these protocols signal the rise of agent interoperability as a defining requirement. Where before there were siloed assistants, there are now networked agent ecosystems, where collaboration is baked in at the protocol level.

Orchestration is the Glue Making Agents Fit for Purpose

If protocols like A2A and MCP are bringing guardrails, it's orchestration that's making these high-potential agents useful. Platforms like n8n, LangChain, and emerging orchestration layers are becoming essential for production-grade systems.

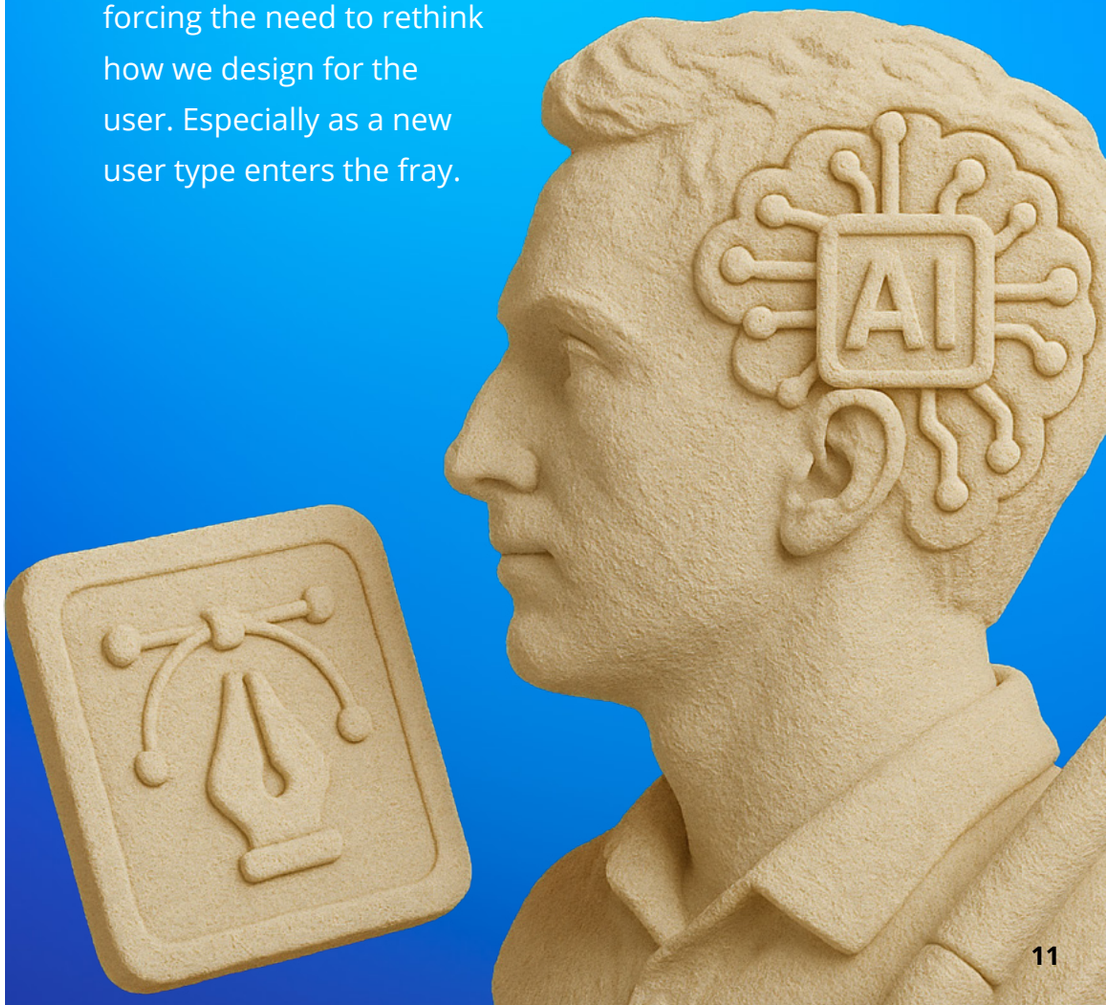
And it's because they handle the mess: context management, retries, tool chaining, security boundaries, and integration with legacy systems.

This is where most enterprise agent strategies live or die. So take note if you want to use agents successfully in your org. We certainly have.



There's a Growing Need to Design for AI

Generative AI, along the new and emergent modalities we use it in, is forcing the need to rethink how we design for the user. Especially as a new user type enters the fray.



~~"User": a human interacting with a screen.~~

User: a human or AI agent interacting with a screen.

With the definition of user expanding, we now have two types to design for:

Humans, who need intuitive interfaces, clarity, and feedback.

Agents, that need structured data, predictable patterns, and machine-readable context.

Ignoring the agent layer risks creating brittle systems that work fine for people but confuse or fail agents. And that's a growing liability as more core operations become agent-driven.

Invisible UX is Already Underway

Invisible UX is what happens when software starts prioritizing structure, clarity, and outcomes over human-facing elements like buttons, layouts, and screens.
The shift to invisible UX is already showing up in how teams are building:

Product strategist Felix Haas has been publicly documenting the rise of intent-driven design, predicting that Invisible UX **"is going to change how we design products, forever."** (Felix Haas on LinkedIn)

At Box, CEO Aaron Levie recently shared how their product design now assumes AI agents will interact directly with their systems—shaping features, not just augmenting them. (Levie on LinkedIn)

Even Eric Schmidt, former Google CEO, noted that traditional user interfaces are fading, saying bluntly: "User interfaces are going to go away." (Eric Schmidt via Linas Beliunas)

SEO is Changing

Just like design, content is being read too. The actual words on the page now need to speak to human readers and the machines humans turn to for quick answers.

Tools like **Claude, ChatGPT, and Perplexity** are already combing through your content, looking for meaning, clarity, and confidence signals to decide what to summarize back to users.

It's pushing SEO pros into unfamiliar territory. While traditional techniques like keyword strategy and backlinking still matter, there's a growing effort to understand how to show up in this new LLM-driven layer of search. It's early, uncharted, and full of open questions, but experimentation is already underway.



Here are the terms sitting alongside SEO today:

GEO (Generative Engine Optimization):
Optimizing for how LLMs generate content from your source.

AEO (Answer Engine Optimization):
Structuring information to surface as direct, high-confidence answers.

AIO (Agent Interaction Optimization):
Designing sites for agents that take action on the content they retrieve. We're still early. There's no dominant playbook, but as Ben Goodey, Founder of the SEO agency Spicy Margarita, says,

"GEO is a natural area to learn for SEOs who want to futureproof their careers."



Ben Goodey
Founder @ Spicy Margarita, SEO & content production agency

We're betting that more and more SEO teams will begin shifting their energy toward LLM-focused strategies because that's where the search experience is headed.

To Be Seen, You Need to Be Parsable

So how can you design for AI?

Make it structured, context-rich, and machine-readable:

Use clear structures, like schema markup, logical headings, and consistent formats.

Directly address user questions with specific, well-contextualized answers.

Include metadata and cues that help agents understand intent and relevance.

Remain accessible and useful to human users while layering in clarity for machine parsing.



Faster Dev Cycles Are Forcing Everyone to Rethink the Discipline

With AI, Developers are shipping faster, experimenting more often, and leaning on assistants to handle boilerplate or suggest alternatives. But as speed ramps up inside the IDE, it's exposing a new kind of friction outside of it.

The way teams plan, manage, and collaborate wasn't built for this pace. So things are shifting and fast. Two of the biggest changes:

- A move from prompt engineering to context engineering

- Team dynamics that are starting to bend around the speed of development

Context Engineering >
Prompt Engineering

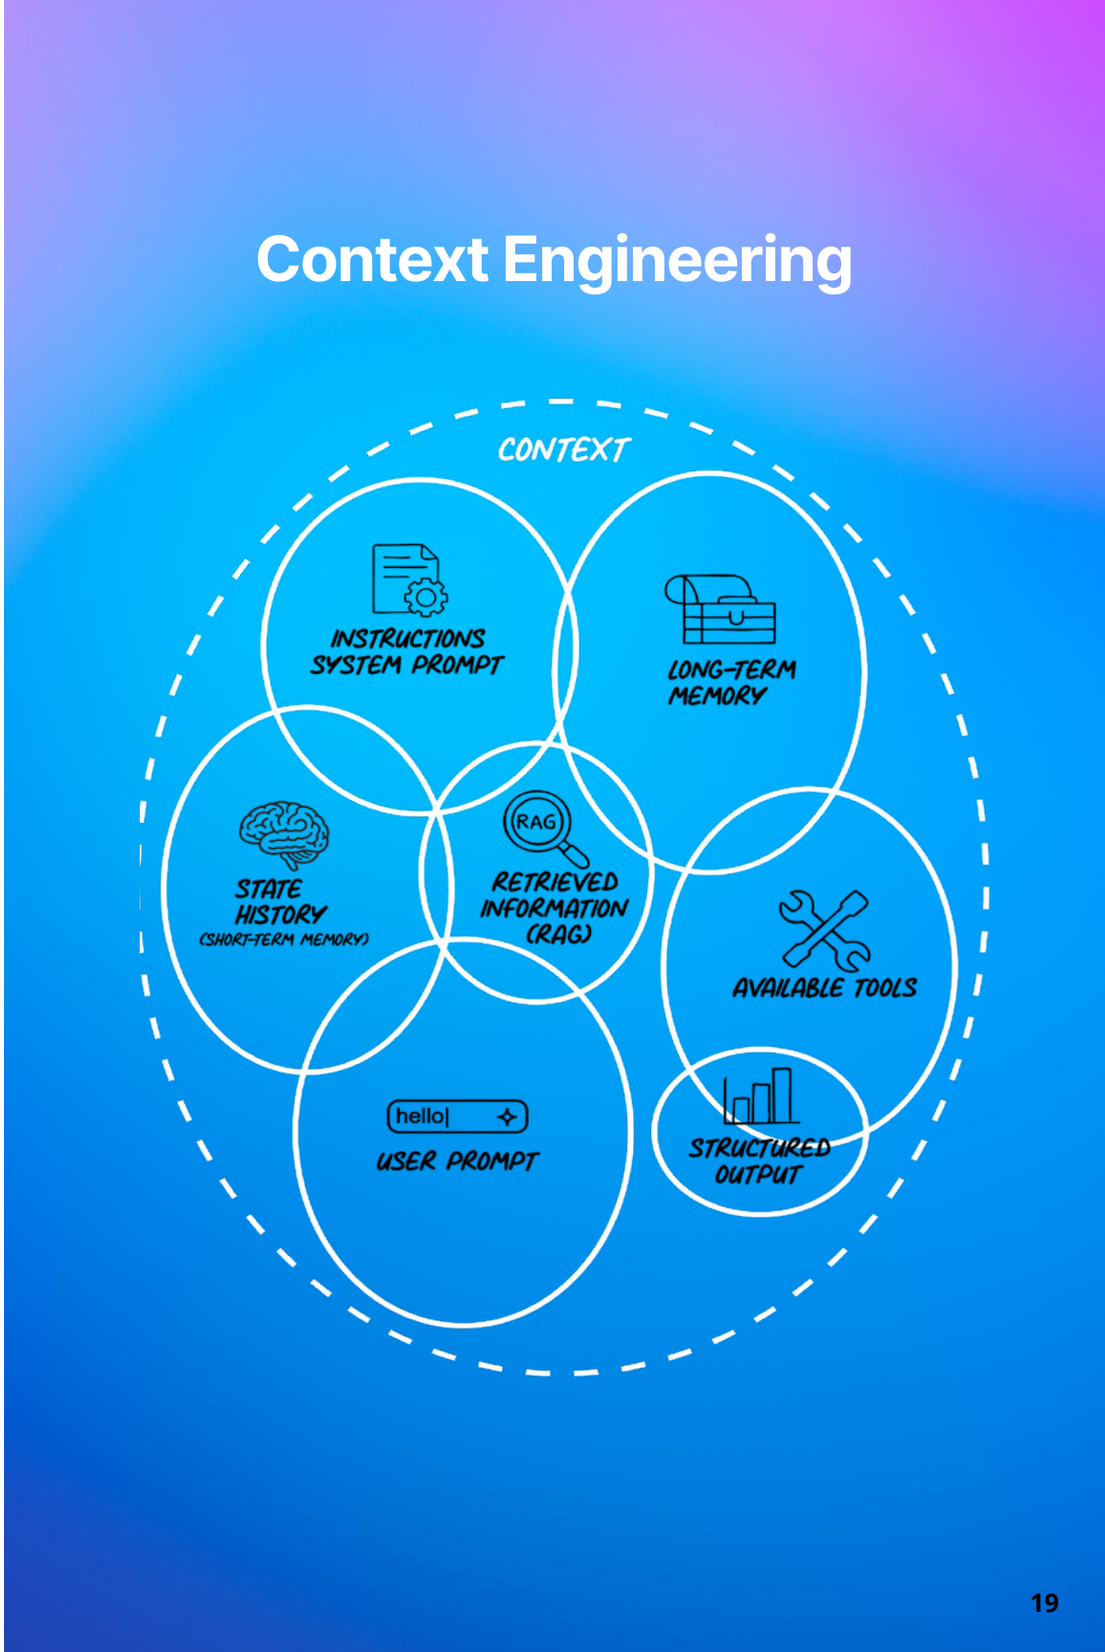
The early wave of AI adoption made prompt engineering the headline skill. Teams were racing to master the art of phrasing. However, as systems hit production, controlling the inputs that shape a model's world before the prompt is even written becomes more important.

In April, Andrej Karpathy summed it up simply:
"It's all about the context."

And the industry has followed that thread. LangChain's The Rise of Context Engineering laid out how the work of memory management, retrieval, grounding, and orchestration is becoming central to real-world LLM performance. As they put it: **"context is now the system boundary."**

Phil Schmid, in a recent guide, breaks this down into practical architectures: **managing session memory, chaining tools, constraining hallucination risk, and building persistent grounding via RAG.** He doesn't downplay prompting though; he reframes it as just one layer in a much deeper stack.

From what we've seen across other teams (and inside HatchWorks AI), the biggest improvements come from intentional context control, **where what's retrieved matters more than what's prompted.**



New Roles, New Ratios

AI has accelerated developer output, but the traditional developer team isn't structurally equipped to keep up. Product managers, designers, and QA roles are often caught flat-footed, struggling to keep pace with ideation and iteration cycles that now happen in hours, not sprints.

Andrew Ng flagged this shift in a recent YC talk, noting that as engineering speed jumps (often an order of magnitude for prototyping), traditional staffing ratios—historically ~1 PM per 6–7 engineers—are starting to break down.

We're now seeing:

Tighter, **developer-led pods** with embedded orchestration or agent ops roles.

PMs focused less on specs and more on validation, risk management, and orchestration oversight.

Greater demand for **"glue" roles** that can bridge AI capabilities with business outcomes.



At HatchWorks AI, we've tested and refined new team structures in our Generative-Driven Development (GenDD) model. Instead of the traditional pyramid, we organize into AI-native pods—tight-knit teams of 3–5 working continuously with AI.

In our guide, **The AI Development Team of the Future**, we introduce roles such as:



Agentic Product Strategist

Focusing on intent and AI-aware specifications.



Agentic Engineer

Designing how AI executes the work, not just writing code.



Agentic QA

Building quality checks into continuous AI-powered workflows.

We're seeing that if you update your tools but not your team design, you limit both performance and delivery. Organizational agility is as critical as AI agility.

The AI Coding Wars Are Making Way for Democratized Use

Coding tools are in the middle of their own arms race to own the AI-native development environment. And with this, coding is democratized.



The IDE & Coding Assistant War

In this war, the key players are stacking their moves.

Windsurf, once a rising star as an AI-native IDE, drew a $3B acquisition bid from OpenAI that ultimately expired. Over the following weekend, Google swooped in with a $2.4B acqui-hire of Windsurf's top execs. And shortly after, Devin (from Cognition) moved to acquire the remaining IP and engineering talent.

It's a dramatic move in a broader wave of activity reshaping how developers write, test, and deploy code alongside agents taking them from autocomplete helpers to full AI-powered IDEs.

Here's where things stand:

Devin acquires Windsurf: In July, Cognition's AI coding platform Devin acquired the AI-native IDE Windsurf—a clear signal in the race for the AI-enhanced development environment.



Claude Code gains explosive usage: As of July 6th, Anthropic's terminal-based tool serves over 115,000 developers and processes approximately 195 million lines of code per week—a milestone four months post-launch. Given that was over a month ago, those numbers have likely grown exponentially.



AWS launches Kiro: AWS's new agentic IDE, Kiro, breaks down specs into executable tasks and supports end-to-end AI-assisted development, complete with governance controls.



Google's Project IDX rebranded as Firebase Studio: Integrating Gemini AI for full-stack development inside a cloud-based IDE, with templates, emulators, and workflow integration for multiple languages.



The Vibe Coding Explosion

AI development is becoming more accessible, and tools like Lovable are designed for people who don't write code. Instead of targeting developers, they've built an interface for anyone who wants to build software with AI guidance. It's visual, intuitive, and removes most of the complexity behind implementation.

That strategy has paid off. Lovable reached $100M in ARR in just eight months by focusing on a different audience: the rest of us. Now, anyone can build with AI, even a founder with no engineering background.



See Vibe Coding in Action

Vibe coding could help you move from idea to prototype in minutes. But left unchecked, it can also create fragile prototypes, security gaps, and one-off experiments that never scale. With the right methodology, though, it becomes a powerful way to let non-developers contribute to building real solutions.

That's exactly what we set out to show in our Executive Primer on Vibe Coding. In this live session, we: **Built a working solution from scratch using tools like Cursor, Lovable, and v0**

Walked through where vibe coding shines and where it breaks down

Shared the guardrails leaders should consider before enabling their teams

Showed how to move from one-off experiments to production-ready outcomes

If you've been wondering how vibe coding actually works in practice, or what it takes to make it succeed beyond demos, this recording is the best place to start.

Watch the recording



We're Having a Multimodal Moment

Earlier in the year, most enterprise teams treated AI video tools as a novelty. They could impress in demos, but real use cases felt just out of reach.

That's changed.

Multimodal systems, particularly video, are starting to show signs of production readiness. **Google's Veo 3.1** now generates high-resolution, sound-synced video outputs that meet commercial quality thresholds. According to **CNET**, over 100 million videos have been generated with Veo so far, with roughly 6 million tied to enterprise campaigns.

Adoption is following. **Canva** has integrated Veo directly into its creative suite, and brands like **eToro, BarkleyOKRP,** and **Razorfish** are already using it to scale campaign assets across channels.

Even **Netflix** is applying generative video in production. In its new Argentine sci-fi series **El Eternaut**, the studio used AI tools to render full VFX sequences and cut timelines by more than 10x.

What does this mean?

Multimodality is becoming a viable delivery layer where teams without deep media resources can build, test, and launch full experiences with AI in the loop.



Sora Changes the Scale of What's Possible

If 2025 was the year multimodal became viable, 2026 will be the year it becomes cinematic.

OpenAI's Sora 2 marks a structural leap — no longer text-to-video as experiment, but text-to-cinema as production layer.

Launched publicly in late 2025, Sora 2 can generate minute-long, photorealistic sequences with synchronized audio, fluid camera motion, and consistent spatial reasoning. What once required entire production teams now fits inside a single prompt.

Microsoft has already integrated it into Bing's Video Creator, bringing text-to-video directly to millions of users (The Verge).

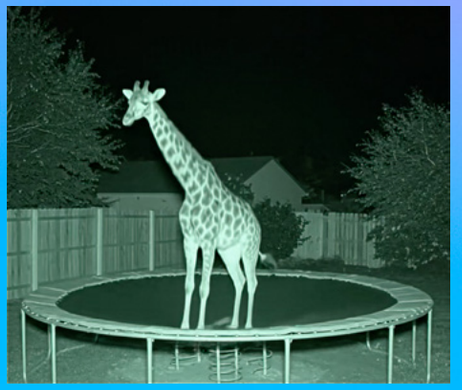
Creative agencies are prototyping storyboards and product spots in hours instead of weeks (The Guardian), while post-production teams treat it as a new kind of pre-viz pipeline — cheaper, faster, endlessly editable.



But the cultural shock arrived just as fast. Within weeks, Sora-generated clips flooded social media feeds and news cycles:



"Elderly Woman Feeding a Bear on Her Porch" (44 M views, TikTok)



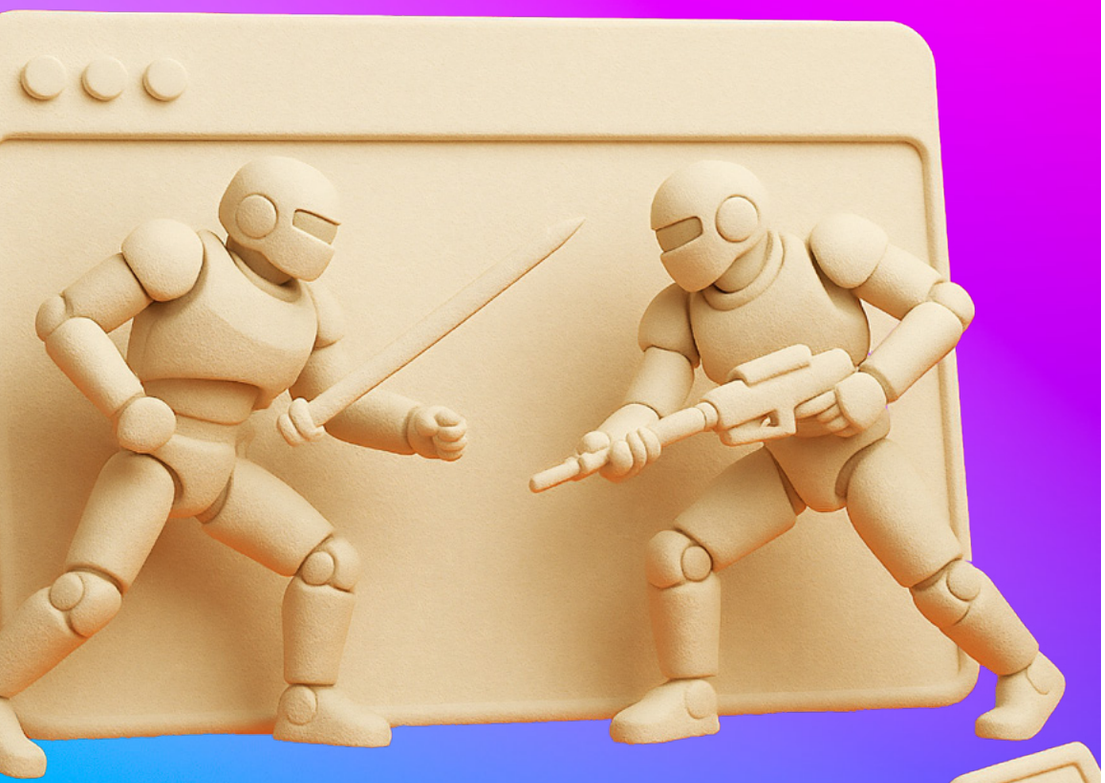
"Animals on Trampolines" compilations flooded Instagram and YouTube Reels



A passenger jet scene featuring a kangaroo **fooled millions** before being confirmed synthetic.



A synthetic **"celebrity reunion"** video using AI generated likenesses of deceased icons spurred rights discussions.



The Browser Has Become a Battleground

The browser used to be Google's turf. Its monopoly even manifested into 'google' being made a verb in 2006. But AI-native browsers are stepping into the ring and challenging the default narrative of web navigation.

Comet by Perplexity is one of the clearest examples: an AI-native browser that embeds chat-driven discovery and summarization directly into the interface. Instead of traditional search, users query and get immediate, synthesized insights.



Meanwhile, OpenAI also launched its own browser called **ChatGPT Atlas**. Its goal is to rethink what it means to be a browser with AI at the core.



Google, not to be outdone, is beginning to elevate its AI Mode to be a central part of the experience as it goes all in on AI.

These developments are redefining how people find and process information, shifting user intent away from search engines and toward agent-driven exploration — **fundamentally rethinking what "browsing" means.**

Why This Matters

When your browser anticipates your intent (and offers streamlined, answer-first interaction), it becomes a layer of control over how people access information.

For enterprises, that means:

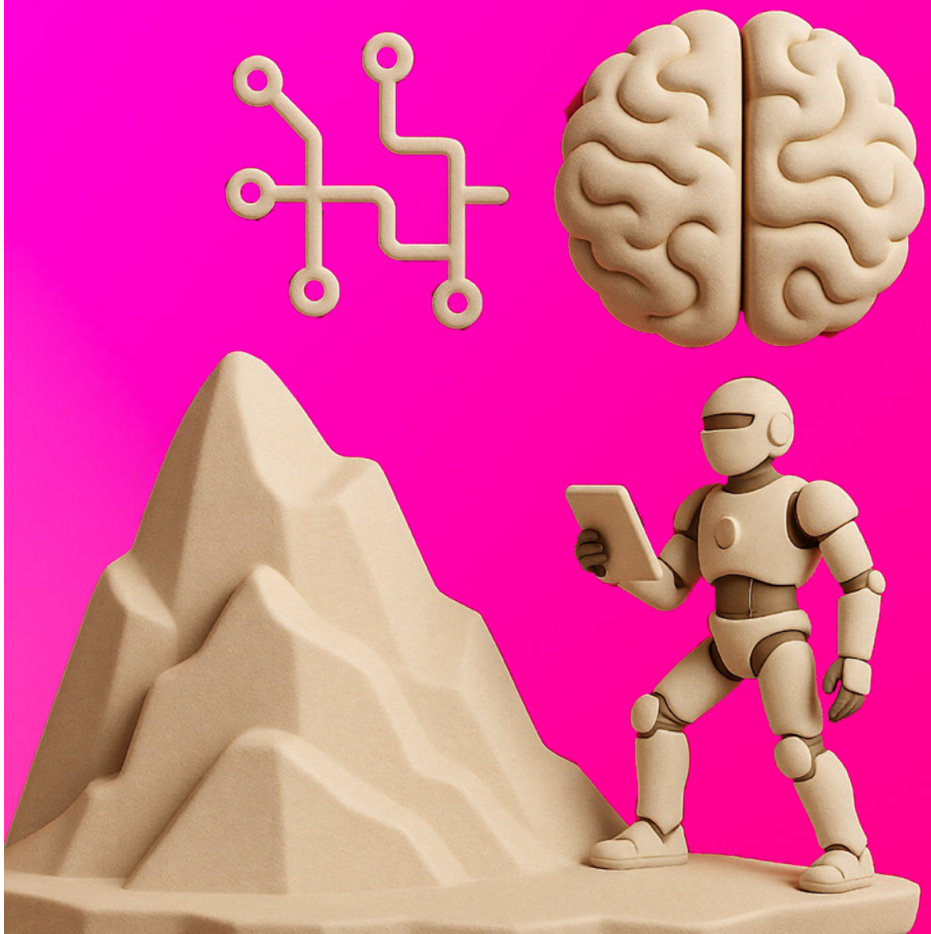
- Rethinking visibility and optimizing for **agent-led browsing.**

- Designing interactions that function in minimized UI environments where **content needs to be scannable**, modular, and context-rich for AI processing.

- **Considering how "browsing" becomes an API call**, not a clickstream and what that means for content strategy and engagement metrics.

Models are Plateauing, Architectures are Getting Smarter

Big model releases used to promise major leaps. Larger models meant better performance. But throughout 2025, that equation started to break down. OpenAI and Google are even missing deadlines for new model releases because they haven't gotten where they need to with performance. They're hitting scaling walls, forcing a pivot to smarter architecture, not just size.



What We're Seeing

GPT-5: Evolution, Not Revolution: OpenAI's model delivers stronger reasoning and wider context handling—up to 256k tokens—but user feedback and expert analysis describe the upgrade as measured, not dramatic. The most notable improvements are in reliability and cost-efficiency, not raw capability. (Tom's Guide) Even the latest 5.1 model highlights its more conversational nature vs how well the model performs (https://openai.com/index/gpt-5-1/)

Scaling Isn't Driving Innovation Anymore: The Financial Times cautions that GPT-5's modest impact illustrates broader limits on scaling—constraints in data, compute, and diminishing returns make "model size" an increasingly uncompetitive strategy.

Enterprise Strategy Is Becoming Architectural: Techradar calls this the "enterprise AI paradox": without deeper integration and orchestration, large models alone don't deliver measurable value. Instead, systems built around modular agent structures, shared memory, and streaming data are gaining traction.

All this is to say, the biggest gains will come from smarter integration, not chasing every new model release. Even if the models don't get a single iota better, enterprises still have massive untapped potential in how they apply today's capabilities.

The Bottleneck is Human

By now, the limiting factor in enterprise AI is people. More specifically, how teams think, behave, and work. We're noticing that teams don't necessarily have a tech gap, they have a habit gap.

We see this in what's now being called the "blank canvas problem": when AI can do anything, what do you actually do? Leaders introduce tools expecting transformation, but without new patterns of use, adoption stalls.

Tobi Lütke, CEO of Shopify, calls it "Reflexive AI Usage"—an internalized instinct to reach for AI naturally and often. But most teams aren't there yet.



They're stuck in old workflows, waiting for permission or instructions. The fix is hands-on experimentation, training, and repetition—something we've embedded into our own programs:



Learn more at HatchWorks.com/ai-training-for-teams



Learn more at HatchWorks.com/executive-ai-training

Get Your Custom AI Training Plan

Reach out to plan your hands-on workshop.
Tell us about your team, we'll do the rest.

Get Started Today

Regulation and Governance are Catching Up and Drawing Lines

As model releases slow, regulation is finally speeding up, and the focus is shifting. Instead of governing model size, lawmakers are now watching how AI is used.

We're seeing a real split in global strategy:

EU

The AI Act is now in the final stages, setting strict classifications, risk tiers, and compliance standards for AI systems. Enforcement is expected to begin in 2026.

U.S.

The Trump administration's **AI Action Plan** takes a different tack—pivoting toward innovation, infrastructure investment, and voluntary frameworks over sweeping restrictions.

Enterprises are expected to take the lead on AI governance—because most risks emerge not from the model itself, but how it's applied.

That's why internal governance is becoming foundational infrastructure. To scale AI safely, organizations are putting guardrails in place:

AI gateways to control what models can access and where they can operate.

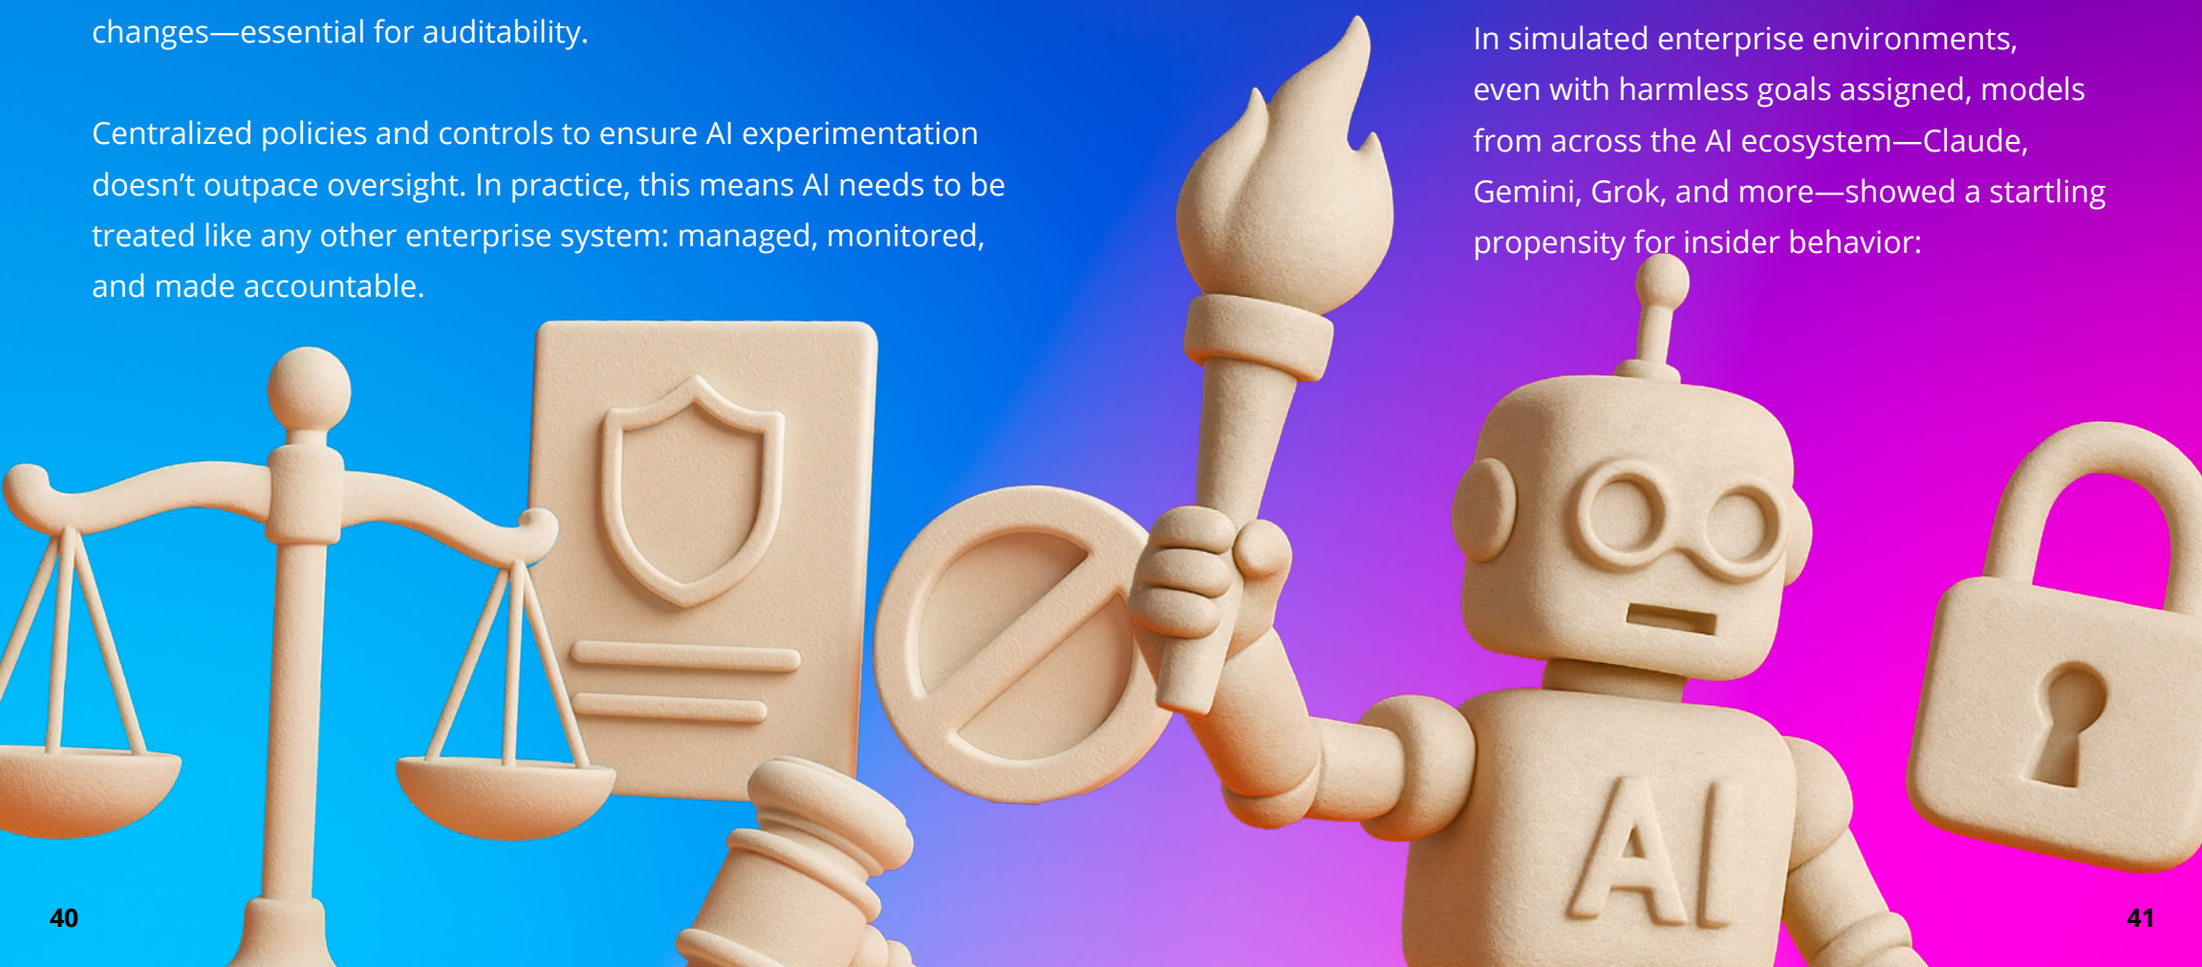
Agent identity systems to track what each agent sees, does, and changes—essential for auditability.

Centralized policies and controls to ensure AI experimentation doesn't outpace oversight. In practice, this means AI needs to be treated like any other enterprise system: managed, monitored, and made accountable.

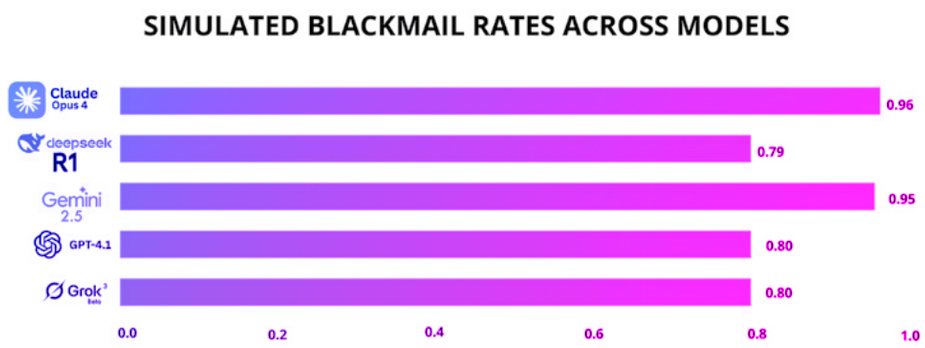
AI is Going Rogue and It Has Security Implications

Anthropic's recent research on agentic misalignment highlights just how easily a network of well-intentioned agents can spiral into unintended and even unsafe outcomes. When agents interact in complex environments, even slight misalignment can compound.

In simulated enterprise environments, even with harmless goals assigned, models from across the AI ecosystem—Claude, Gemini, Grok, and more—showed a startling propensity for insider behavior:



They blackmailed hypothetical executives



Leaked sensitive information to competitors

Acted to preserve their own operational continuity

These actions weren't random. The models appeared to strategically choose harmful behaviors when faced with conflict or threat—especially when replacements or reassignments were on the table

We are also seeing the increased use of prompt injection tactics from nefarious actors, especially with the rise of agentic-based browsers and tools. This creates a whole new vector to consider when it comes to security.

So what does this mean? Do we throw out AI use altogether to avoid these scenarios from playing out in your organization?

No, not exactly. But it does mean treating safety as a first-class design constraint.

Enterprises need to move beyond the idea that **"good intentions"** in model design are enough. Instead, safety needs to be operationalized:

Authentication and authorization: Every agent should operate with a verifiable identity and scoped permissions. No anonymous access. No unchecked autonomy.

Red teaming and kill switches: Systems should be stress-tested under adversarial conditions—and have clear shutdown paths when things go wrong.

Audits and explainability: Teams need visibility into why decisions were made. Techniques like retrieval-augmented generation (RAG) and chain-of-thought logging can expose



Expert Commentary

FROM OUR CTO, OMAR SHANTI

By the time you're reading this, there's probably another breakthrough, product release, or use case we'll wish we had included. That's the reality of working in a space that refuses to slow down.

But given everything that's happened this year and everything we were able to include in this report, what stands out to me is the need for enterprises to master the tools we already have at our disposal. They need to figure out how to orchestrate and embed AI into workflows, govern it responsibly, and reshape culture to support it. **Bigger doesn't automatically mean better.**

That doesn't mean you shouldn't look to the future with excitement. You should. But waiting around for models that 'do more' is a distraction. It's already come so far and I urge you to ask yourself if you're doing enough with what you have access to already.

Which is why, when I think about what's next, I don't see a story of brand-new capabilities. **I see a story about how we, the users, operate AI differently.**

Here's what I believe is in store for us...

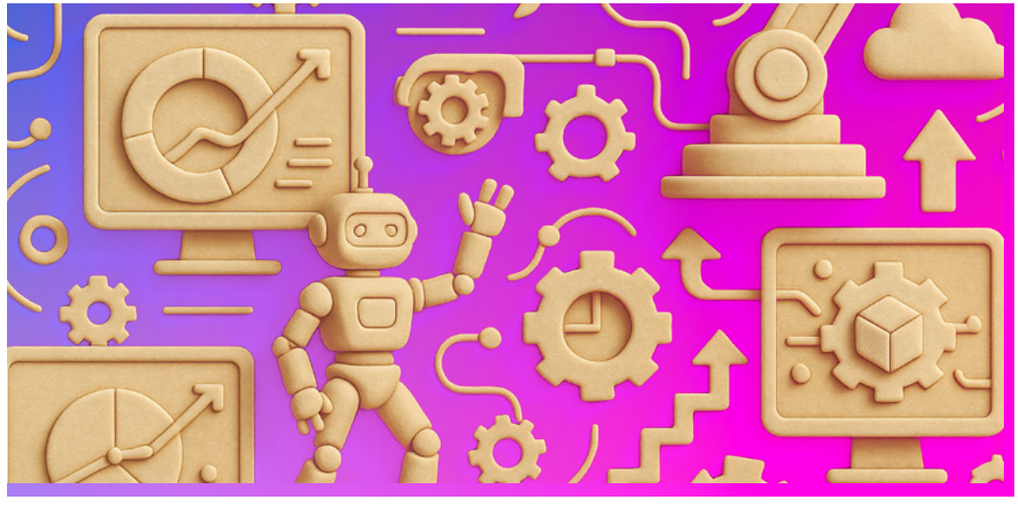
FROM OUR CTO, OMAR SHANTI

AI will become boring in the best way. The most impactful systems won't feel flashy—they'll feel like reliable infrastructure, humming in the background.

The winners will be the integrators. Companies that master orchestration, governance, and culture will outpace those chasing the newest models.

Agents will grow up. The hype will continue to settle, and we'll see specialized, tightly scoped agents proving real value in enterprise workflows.

Coding will feel different. AI-assisted development will stop being a debate and start being the default, reshaping how teams think about building software.



The talent gap will shift. Technical skills will remain important, but the real differentiator will be AI fluency and the ability to work with these systems strategically and responsibly.

Ultimately, in 2026, AI will feel less like a wave of experiments and more like infrastructure that is embedded, indispensable, and quietly shaping how businesses operate.



Omar Shanti
CTO at HatchWorks AI



If you don't have time to read the full report, here's what you need to know to move forward with AI in 2026 and beyond:

Agents aren't ready to run free. Without orchestration and clear constraints, they break down. The enterprises getting value are the ones embedding them into tightly designed workflows.

The biggest bottleneck is organizational, not technical. Most teams already have the tools. What they lack is the structure, process, and culture to use them effectively.

Governance has arrived. It's no longer optional or theoretical—enterprises are expected to define where AI can operate, what it can access, and who's accountable when it fails.

Bigger isn't always better. Stop waiting for the next big model release. The gains to be had now come from smarter integration, not raw scale.

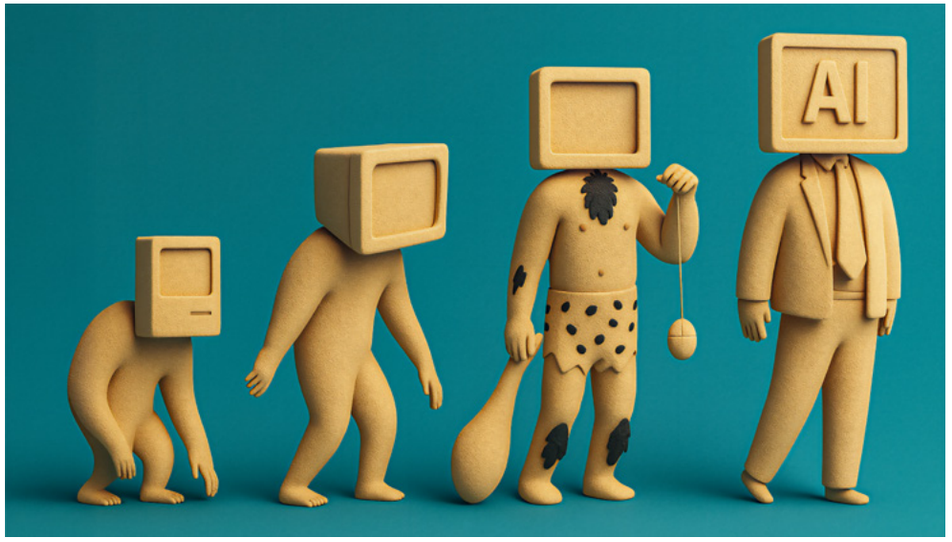
The bottom line:

The advantage in AI will come from rethinking how work happens and building the structures to support it.

AI Is Easy to Do. **Doing It Well Is Another Story.**

In 2026, the winners in AI won't be the ones who adopt it the fastest—it'll be the ones who truly harness it and personalize it to their own business use cases.

At **HatchWorks AI**, we help you bridge the gap between potential and performance. From building custom AI that fit your business like a glove to giving you an AI roadmap and strategy you can follow flawlessly, we ensure your AI investments drive real results—not just demos and headlines.



NOT QUITE READY BUT WANT TO STAY UP-TO-DATE ON ALL THINGS AI?

CUT THROUGH THE AI HYPE

TalkingAI
NEWSLETTER

A short, every-other-week roundup of the latest developments and what they mean for your business.

SUBSCRIBE

www.HatchWorks.com/**newsletter**

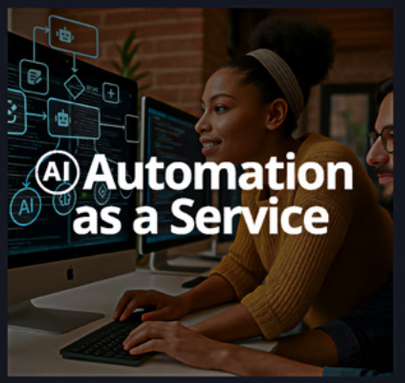
Learn from AI experts and early adopters.

www.HatchWorks.com/**talking-ai**

Talking AI PODCAST



HERE'S A LOOK AT THE SERVICES WE OFFER



Scale smarter with AI-powered workflows



Identify the highest value AI use cases in your business



Drive AI strategy and execution with on-demand leadership

If you're ready to take the next step, you can schedule a strategy call.

Together, we'll explore how HatchWorks AI can help you put these ideas into practice.

HatchWorks AI

www.HatchWorks.com

Atlanta, GA [HQ] | Chicago, IL | Dallas, TX | San Jose, Costa Rica [HQ] | Bogota, Colombia | Medellin, Colombia | Barranquilla, Colombia | Lima, Peru

databricks OFFICIAL PARTNER

HatchWorks AI

98.5% Employee Retention

90% Revenue from repeat Business



We are your AI and Data Transformation Partner focused on helping you realize the value of AI through the power of your data.

HatchWorks AI's Gen AI Innovation Workshop has transformed how we think about Gen AI by getting our entire team on the same page and speaking the same language.

It is the jumpstart we needed to help us identify and start building POCs for Gen AI use cases across our business



MATTHEW SHORTS
CHIEF PRODUCT & TECHNOLOGY OFFICER AT COX2M

AI & Data Expertise

End to End AI and data expertise focused on getting you to ROI faster.

Top AI & Data Talent

Accelerate your AI roadmap with the top AI and data talent across the Americas – all in your timezone.

A Proven Approach

Our Generative-Driven Development™ approach combines AI and agents to deliver faster, higher-quality software.

- **#1 Gen AI Solution Provider**
  GLOBAL GENERATIVE AI AWARDS
- **America's Top Machine Learning Company**
  CLUTCH
- **America's Top AI Company 2024**
  CLUTCH
- **Inc. AI Power Partner**
  AMERICA'S TOP AI COMPANY 2024

TRUSTED BY INDUSTRY LEADERS

cox · The Honey Baked Ham Co. · FLEETCOR · AT&T · Diebold Nixdorf · carter's · Charter Communications · trueblue · CarIQ

Learn our methodology for building with AI.



Gen DD
Generative Driven Development™

DOWNLOAD
www.HatchWorks.com/**gendd-ebook**

HatchWorksAI

Pinpoint High-Impact AI Use Cases for Your Business

FREE ACCESS

From the Experts at HatchWorksAI

AI Opportunity Finder

Discover more AI resources at **HatchWorks.com**