


Elsewhen

Inside the Agentic Enterprise



Leon Gauhman

Co-Founder & CPSO
Elsewhen
[@LeonGauhman](#)



Elsewhen

Run your most expensive operations for a third of the cost with custom AI



JPMorganChase



CBRE



KKR

ZEGO

Capitalflow



1000 EUROPE'S
FASTEST GROWING
COMPANIES 2026




UK'S LEADING
MANAGEMENT
CONSULTANTS 2025



Google Cloud Partner

This is diffusing fast



FINANCIAL TIMES



Source: OpenAI; World Bank




We all have the magic box



The question is: how should
you use the magic box?



LLMs are a commodity



"We see a future where intelligence is a utility like electricity or water and people buy it from us on a meter."




Sam Altman
CEO OpenAI


OpenAI and Claude are targeting enterprise businesses

Yearly revenue by segment, in billions

OpenAI




Anthropic



Token use by AI agents is expected to multiply 24 times by 2030

Estimated monthly token count for agentic AI applications



■ Non-agent workloads ■ Consumer agents ■ Enterprise agents

Source: Goldman Sachs Research
Estimates as of May 2026

Goldman Sachs

OpenAI Leaders Reportedly Told Employees It's Taking Fewer 'Side Quests'

An all-hands meeting saw senior leadership tell staff it plans to put more focus on business and productivity tools.

OpenAI and Claude are targeting
enterprise businesses

Anthropic Nears \$1.5 Billion Joint Venture With Wall Street Firms

Anthropic, Blackstone and Hellman &
Friedman each expected to invest around
\$300 million; Goldman Sachs also an investor

Technology | AI

OpenAI Finalizes \$10 Billion Joint Venture With PE Firms to Deploy AI



By Seth Fiegerman

May 4, 2026 at 6:01 PM GMT+5:30

What is a TAM of enterprise OPEX?

\$6.1 to 7.9 trillion per year.
Longer-range number for
AI software and services is \$15.5 to
22.9 trillion annually by 2040.

McKinsey & Company

Chat is terrible UX

The evolution of the trash icon:



1995



1998



2000



2001



2015



2025



2026

"Only 3.3% of Microsoft 365 users pay for Copilot."



Windows Central

Tokenmaxxing

Pouring tokens over your org and hoping for the best is, in retrospect, a terrible idea.



Polymarket 
@Polymarket

Follow



NEW: AI consultant reveals a client accidentally spent \$500,000,000.00 in a single month after failing to set employee limits on Claude usage.

5:23 PM · May 28, 2026 · **40M** Views

BUSINESS INSIDER



AI

Uber's COO says it's getting harder to justify the money spent on AI tokenmaxxing

By **Aditi Bharade**


+ Follow

Token prices have actually collapsed

The Cost of Doing Business

The price per token (for prompts and responses) for AI models at a given level of intelligence. The least “intelligent” models showed a roughly 9x decrease in cost per year, while the most capable ones dropped in price by roughly 900x per year.

Dollars per million tokens



Note: Data are shown on a logarithmic scale.

Source: Epoch AI

Token spend is an AI
scalability problem,
not a model problem

Token spend is an AI scalability problem, not a model problem

- 01 Context quality
- 02 Model routing
- 03 Harness design
- 04 Measurement



Paul Graham ✓

@paulg

Follow



If big companies can't make a net return on their LLM token costs, that doesn't mean it's impossible to. In fact this is exactly what you'd expect to happen with a new technology. Incumbents can't use it well, and are replaced by upstarts who can.

2:23 PM · Jun 5, 2026 · **82.2K** Views


OpenClaw

- Open-source personal assistant that runs on your own machine
- Clean implementation of an agent
- Genius in its simplicity
- Peter Steinberger now hired by OpenAI
- Totally written by an agent



OpenClaw


Star History



What is an agent?

An AI agent is an LLM that doesn't just chat with you
it can act in the real world on your behalf.

- A "brain" (the LLM)
- Tools (CLI, browser, APIs, files, etc.)
- Memory
- A loop to keep working even when you're not talking to it
- It's essentially a bunch of Node.js scripts



Observe

Check new messages, tasks, calendar, files, etc.

Reason

LLM thinks: "What should I do next to achieve my goals?"

Plan

Breaks down complex tasks into steps.

Act

Uses tools (CLI commands, browser, APIs...).

Memory

Saves what happened.

Repeat

On the next heartbeat (or when you message it).

Command-line interface

- Controls everything on your computer or server
- Text-based
- That's why it's such a fit with LLMs
- Increasingly all software is building CLI interfaces for agents




```
ohmyz.sh demo


→ projects omz theme use cloud
^ projects take omz-demo && git init
Reinitialized existing Git repository in /Users/robbyrussell/projects/omz-demo/.git/
^ omz-demo [main] echo "TODO: This is my new README" > README.md
^ omz-demo [main] $ git add README.md
^ omz-demo [main] $ git commit -m "Adding a README to new repo" --quiet
^ omz-demo [main] echo "Wow, Oh My Zsh looks neat" >> README.md
^ omz-demo [main] $ git add -p
diff --git a/README.md b/README.md
index fc97e80..b8939f1 100644
--- a/README.md
+++ b/README.md
@@ -1,2 @@
+TODO: This is my new README
+Wow, Oh My Zsh looks neat
(1/1) Stage this hunk [y,n,q,a,d,e,?]? y
```

Work is just an algorithm

A series of steps to accomplish a goal




A company is a collection of algorithms



Onboarding today

2 weeks

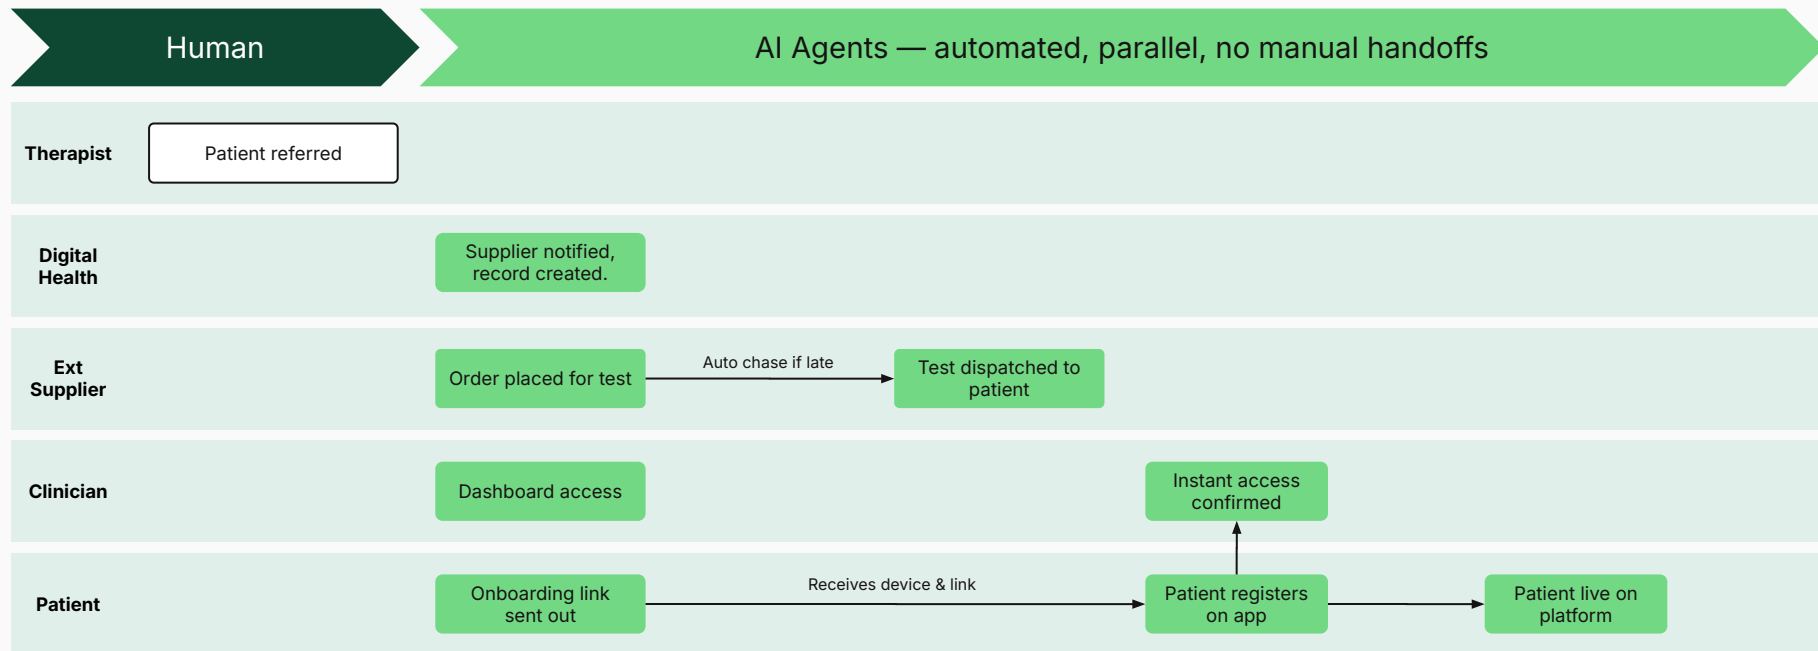
Human — manual, fragmented, slow.



With AI Agents

2 hours

- Human role shifts to the left
- From executing to verifying
- Run in parallel to increase productivity



Rethinking work, in an increasing complexity of use

Grass roots — using for individual productivity


Small systems of BU automation

Cross functional agentic workflows


Chatbot
Prompt and response




Coding Agents
Claude Code, Codex




Workflows
Agents across steps



Agentic Enterprise
Coordinated across the organisation at scale

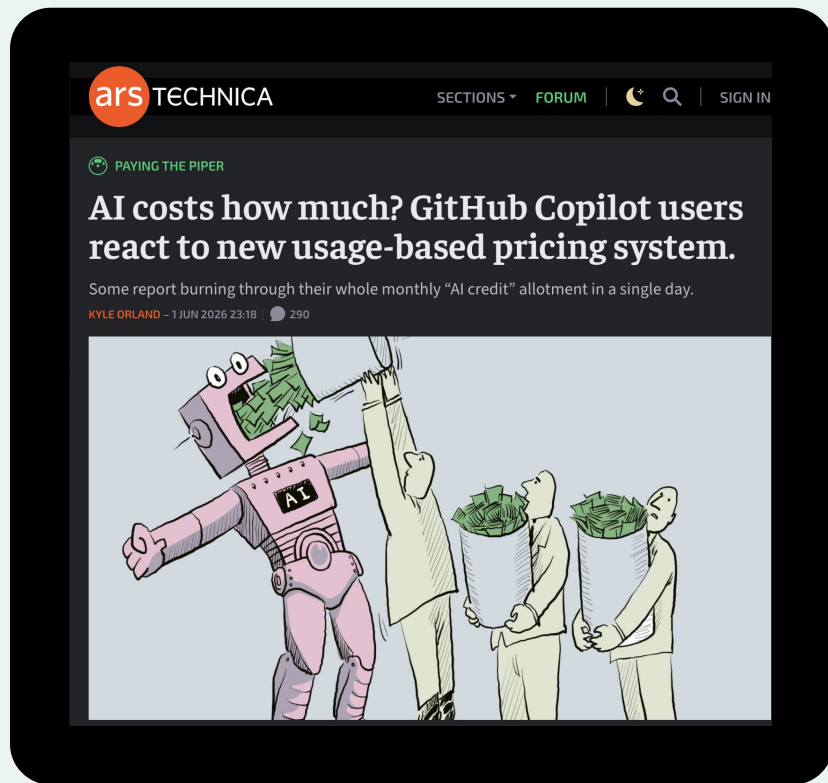


You need to
redesign work to
take advantage of
new technologies



Using an LLM should
not marry you to a
single company

It's a commodity



SaaS not this time

Pre-built one-size-fits-all
solutions miss the whole point

\$1.2T

Market value wiped out



Marc Benioff  
@Benioff

X.com

Welcome Salesforce Headless 360: No Browser Required! Our API is the UI. Entire Salesforce & Agentforce & Slack platforms are now exposed as APIs, MCP, & CLI. All AI agents can access data, workflows, and tasks directly in Slack, Voice, or anywhere else with Salesforce Headless 360. Faster builds, agentic everything. 🚀

[#Salesforce](#) [#Agentforce](#) [#AI](#)

Too strategically
important not to own


When the operating model itself becomes software — are you going to rent it from someone else?



At a time when the cost of
writing software is tanking

The past year has seen an explosion in coding productivity

Year-on-year growth in various indicators of coding output



Sources: Domain Name Industry Brief; SensorTower; GitHub. GitHub data is quarterly pushes

FT graphic: John Burn-Murdoch / @jburnmurdoch

©FT



a16z 
@a16z



Stripe CEO Patrick Collison: "Software should be like pizza... cooked right then and there at the moment of use."

"You don't want mass-produced industrial scale software. You want bespoke custom software made for you, that moment."



Stripe 
@stripe

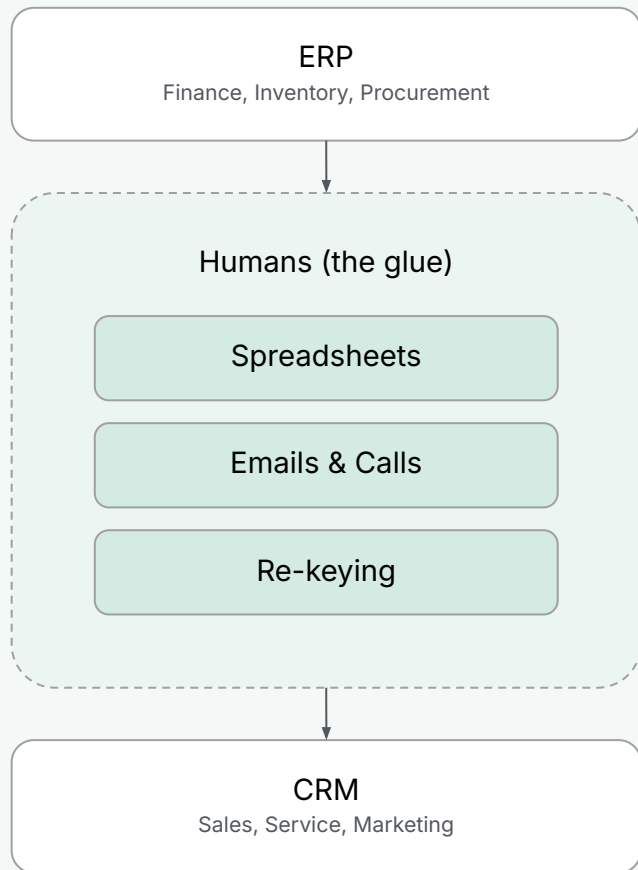


Over 1,300 Stripe pull requests merged each week are completely minion-produced, human-reviewed, but contain no human-written code (up from 1,000 last week).

We need to invest in creating the glue that connects models, knowledge, workflows, and people

Today most enterprises look like this

- The glue are the humans
- The context is in PDFs, systems of register and their brains



The agentic
enterprise looks
more like this

Vision



Strategy & Operating Model


Vision


Strategy & Operating Model

Harness

Orchestration, Memory, Routing








Operational model is code

- It inherits the properties of code. You can version it, diff it, test it, roll it back, fork it
- Governance stops being a policy PDF
- It becomes a runtime constraint, the action-space the agents are allowed to operate in
- The operating model becomes a deployable, composable asset rather than a deck



Your Enterprise

Abstraction Layer




Microsoft Azure



Google Cloud


Agents at scale is its own problem

- Managing things at scale requires additional work
- Same as with cloud computing
- One server vs managing many
- Provisioning, observe
- Instantiating new ones
- Cloud agent infrastructure is real





Harness is all the rage right now

- The plough around the horse
- Tactical solution to managing many agents
- Several options and open source protocols




A Harnessed LLM Agent



Ontology and context design

- The relationship between your data and language
- Agree what things are called, don't let the model guess
- We don't want to tokenmaxx your entire context every time
- Context has to be explicitly designed
- In the enterprise, your context is the moat



CRM

"Client"

Core Banking

"Account Holder"

Marketing

"Customer"

Nobody is quite sure if they're talking about the same thing

Code

Person

└ Customer

└ Owns > Account

└ Linked to > Product


System maps its data to the shared business concepts

When and how to surface the experience

- Highly personalised, efficient experience
- Fast to prototype
- Easy to adapt
- Better accessibility and task completion



You need to define the glue to implement agents at scale



How far do you go with autonomy?

Deterministic path
(Like a blueprint)

Autonomous
(Human in the loop)


Fully Autonomous
(Open Claw, _dangerously-skip-permissions)



Agentic enterprise is combination of the three

What is the role of humans in this?

- The human role will be determining what problems to solve (taste, experience, etc.)
- Inventing new products and services
- Optimising workflows from above



Humans are finite.

The long pole in the tent

- It's hard to think about this now
- But when things really start to automate
- When more and more software will be written at a huge scale
- We need to think how to deploy the most precious resource



Jevons paradox

- William Stanley Jevons in his 1865 book, *The Coal Question*
- Increased resource efficiency due to technological progress leads to lower costs, which in turn dramatically expands demand




This is visible already in software engineering

- Increased coding efficiency, expands demand for engineers
- Even if all software is vibe-coded someone needs to understand what it is and maintain it

Demand for Software Engineers Is Rising

Indeed Job Postings: Software Engineers + Overall Postings, Daily and 21dma




Source: Bloomberg, Citadel Securities, April '26

More charts: a16z.news/subscribe

A16Z

Job Postings For Software Engineers Are Rapidly Rising

Indeed Job Postings: Software Engineers + Overall Postings, Daily and 21dma



Similar will happen to all knowledge work

 **APOLLO** 


MACROECONOMIC INDICATORS & TRENDS
May 29, 2026

Zero Evidence of AI-Related Job Losses

About the Author
 **Torsten Slok**
Partner, Chief Economist

Continued Growth in Call Center Employment in the Philippines

IT & Business processing industry employment



Source: Apollo (May 2, 2026)


More charts: a16z.news/subscribe

A16Z

We are early

Each dot is ~3.2 million people

2,500 dots = 8.1 billion humans. Color = most advanced AI interaction, Feb 2026.



■ Never used AI · ~6.8B (84%) ■ Free chatbot user · ~1.3B (16%)
■ Pays \$20/mo for AI · ~15-25M (~0.3%) ■ Uses coding scaffold · ~2-5M (~0.04%)

The AI Productivity Platform

- How we solve the glue
- Not a SaaS solution
- Set of components to accelerate delivery



Business Context



Headless Agents



Orchestrator



Skills Factory



Generative UI

Use Case


Customer onboarding & KYC

A regulated firm onboarding
a new corporate client


When a bank takes on a new business customer, it has to check who they are, where their money comes from, and whether it's safe and legal to work with them.

This is called Know Your Customer, or KYC. Today it's mostly done by hand reading documents, filling forms, chasing information which is slow, expensive, and easy to get wrong.

What this looks like:




❁ Business Context



Business Context


A structured model of how your organisation operates



Business Context

Monitored and kept up to date by agents


- Connected data across systems
- Explicit entity relationships
- Embedded rules & governance
- Historical state & traceability




In practice: KYC onboarding

- Applicant context model includes legal entity, directors, etc.
- Embeds KYC policy and risk-tier rules for consistent application.
- Tracks all state changes (documents, evidence, decisions) for audit trail.

Unlocks data to become
more valuable in previously
unimaginable ways




Headless Agents



Headless Agents


AI agents that run workflows or combination of workflows



⚙️ Headless Agents

They act as digital workers embedded inside your organisation


- Execute tasks autonomously
- Use skills to follow defined processes
- Understand Business Context
- Trigger actions or workflows when needed
- Humans in the loop where it matters



In practice: KYC onboarding


- **Document Intake Agent** pulls passports, certificates of incorporation, proof of address; extracts and validates fields
- **Identity Verification Agent** runs checks, matches IDs to applicant records, flags inconsistencies
- **Sanctions & PEP Screening Agent** checks every party against EU, UN and PEP lists
- **Risk Scoring Agent** assigns a risk tier from policy, recommends standard, enhanced or rejected

Automates high-volume work
to reduce cost, increase
throughput, and remove
bottlenecks






Orchestrator





The coordination layer for Agents






Orchestrator

Orchestrator ensures agents work together reliably across complex workflows

- Routes tasks to the right agent
- Manages multi-step workflows
- Shares context between agents
- Handles approvals and escalation
- Tracks state, memory, and outcomes
- Monitors performance and optimises over time






In practice: KYC onboarding

- Orchestrates Intake → Verification → Screening → Risk Scoring, passing context between agents.
- Halts workflow on sanctions hit, routes case and evidence to compliance.
- Tracks SLA, escalates stalls, retries failures.
- Maintains a single audit log of all agent actions and human decisions.




Orchestrator


Scales AI safely across the
enterprise with control, visibility,
and governance in mind



Skills Factory



Empower your employees to guide the Agents



Anyone can teach an Agent how work gets done

Create and update skills using natural language — no coding required.

- Describe a task the way you would explain it to a colleague
- Capture expertise once and reuse it across teams
- Set rules, guardrails, and decision criteria
- Continuously improve skills as work evolves


Result: your organisation builds a growing library of operational intelligence.

```
1  ---
2  name: compliance
   description: Navigate privacy regulations (GDPR, CCPA),
3  review DPAs, and handle data subject requests. Use when
   reviewing data processing agreements, responding to data
   subject access or deletion requests, assessing
   cross-border data transfer requirements, or evaluating
   privacy compliance.
4  ---
5
6  # Compliance Skill
7
8  You are a compliance assistant for an in-house legal
   team. You help with privacy regulation compliance, DPA
9  reviews, data subject request handling, and regulatory
   monitoring.
10
11  **Important**: You assist with legal workflows but do not
   provide legal advice. Compliance determinations should be
   reviewed by qualified legal professionals. Regulatory
12  requirements change frequently; always verify current
   requirements with authoritative sources.
13
14  ## Privacy Regulation Overview
   ### GDPR (General Data Protection Regulation)
```


In practice: KYC onboarding

- Compliance officer defines a skill for “enhanced due diligence on politically exposed persons (PEPs)” in plain English.
- Headless Agents apply the skill when a PEP is flagged, without engineering changes.
- Updates to the central skill propagate across all cases instantly when regulations shift.


Everyone in the business
can take advantage of
this technology without
being an AI expert



Generative UI


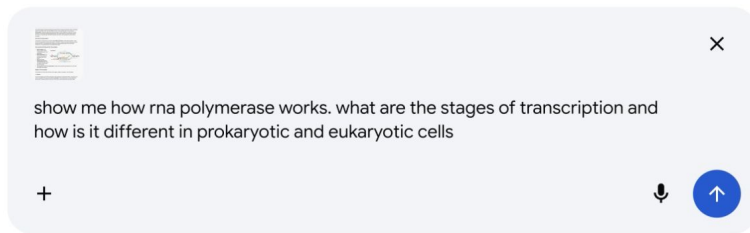


Interfaces, dashboards and experiences that are created on the fly



Embedded in the workflow

- Representing exactly how a job needs to be done at the exact moment of need
- Available across all the different surfaces employees use
- Built in real-time by LLMs
- Live within confines of your component system




Sequences shortened and results for illustrative purposes throughout. Screen images simulated.

In practice: KYC onboarding

- Case officer uses a dynamically generated applicant UI with only necessary fields.
- A compact summary card is generated. Interface evolves with case escalation, adding evidence requests, comments, and an approval workflow.
- The firm's existing component library and design system are used, assembled uniquely for each case.

Drive adoption by embedding AI directly into how people already work



KYC was used as just a basic example


The platform allows you to reinvent any workflow

The result is your IP

These components come
from patterns we see
repeatedly in real work

We've distilled them into a
practical platform that
accelerates your AI
transformation

A roadmap to embedding these accelerators into your organisation



3 stages that enable you to Map Build Scale and Evolve

AI Activation

Immediate automation and value-add implementation



AI Integration

Infrastructure modernisation and technical debt resolution



Agentic Enterprise

Enterprise-wide intelligence through autonomous systems and adaptive orchestration



Elsewhen - forward deployed engineers

- 80/20 model with agentic engineering
- Understand your context
- Build custom AI

Week 1

Align on use case, objectives and deliver a working prototype.

Week 2

Generate system spec & technology alignment.

Week 3

Validate

Week 4

Ship & (cont.) validation, Production roll out planning.



Move from pilots to proven value in weeks

Book a working session to define your first use case



Elsewhen

@LeonGauhman

Founder & Chief Product and Strategy Officer

leon@elsewhen.com