* Booping…

Bring Your Own Agent

We'll Bring the Cluster

Natale Vinto
Technical Director, OpenShift

Cedric Clyburn
Developer Advocate, Red Hat AI



Red Hat

Your Journey to LLM Awesomeness



Prompt

Your Journey to LLM Awesomeness



Prompt



Structured Output

Your Journey to LLM Awesomeness



Prompt



Structured Output



Tools MCP



RAG

Your Journey to LLM Awesomeness



Prompt



Structured Output



Tools MCP



RAG



Patterns

(Evals, LLM-as-a-judge, maker-checker, ReAct)

Your Journey to LLM Awesomeness



Prompt



Structured Output



Tools MCP



RAG



Patterns

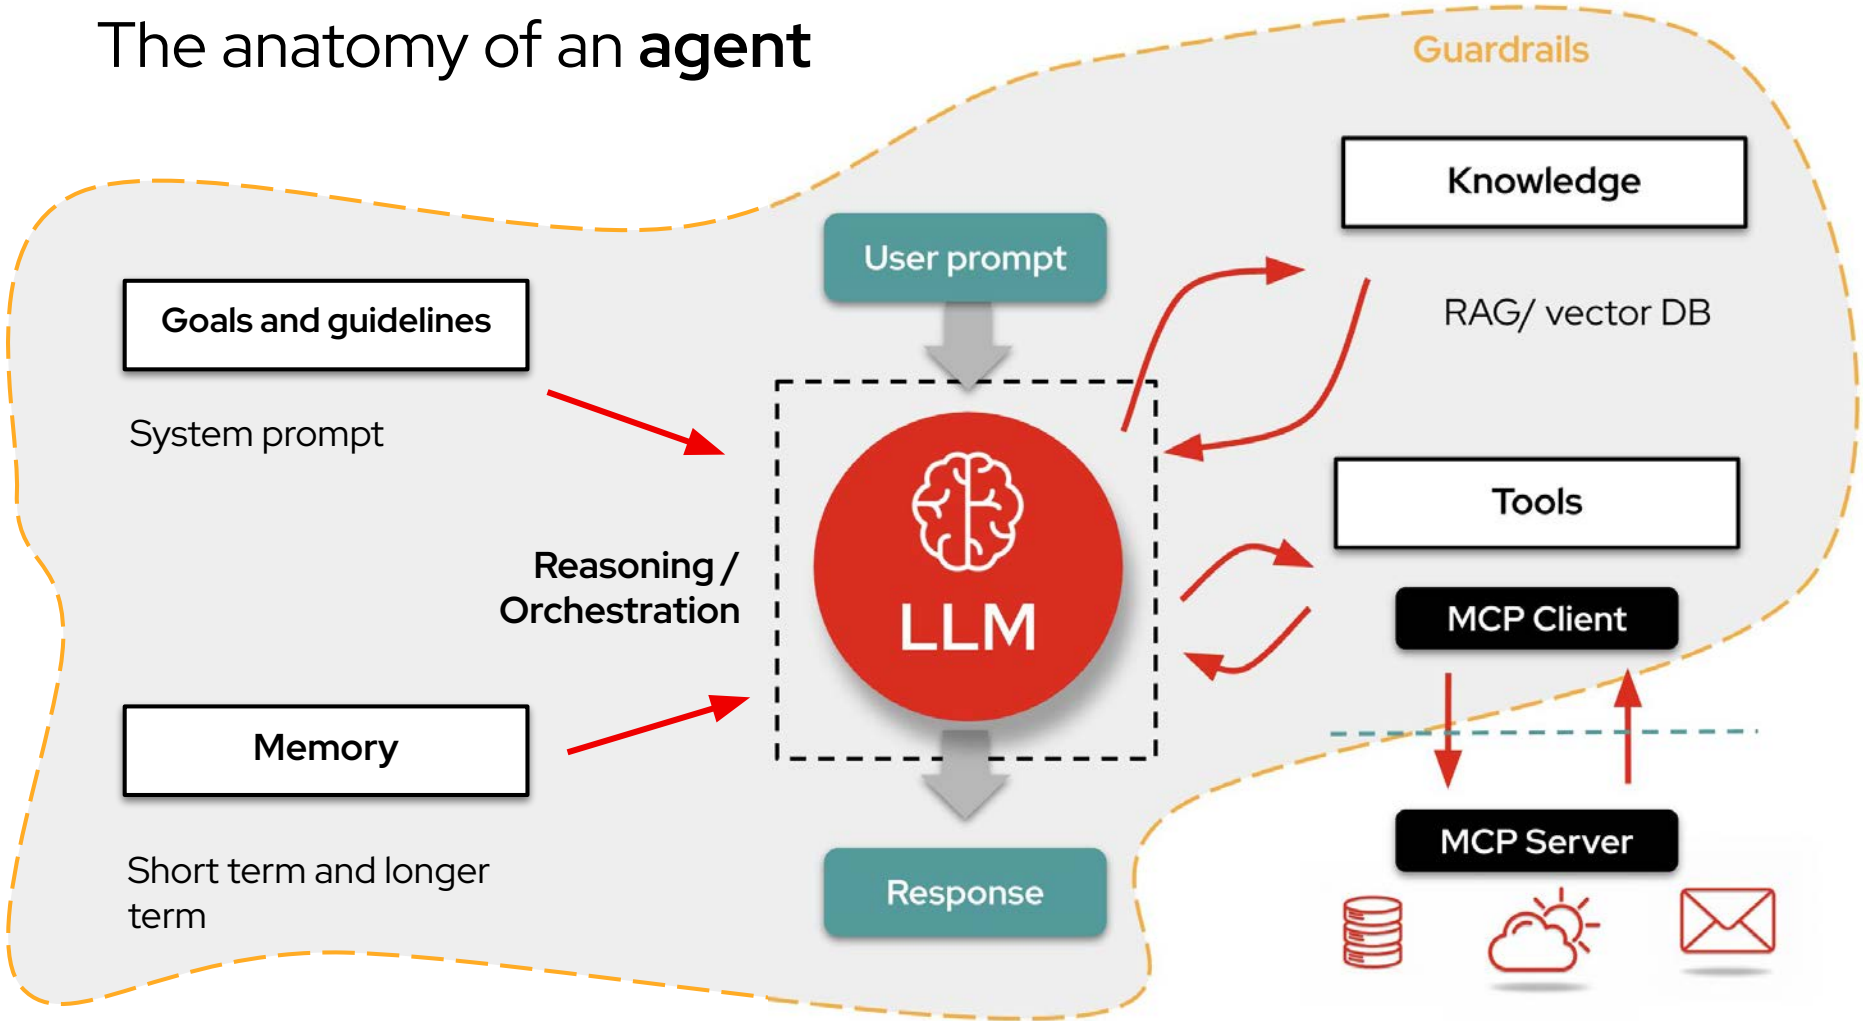
(Evals, LLM-as-a-judge, maker-checker, ReAct)



Agents

(planning, orchestration, retries, error-correction)

The anatomy of an **agent**



Tools
(dynamic data, actions)





"How many ladders do we got inventory?"

User Query

System Prompt
Tool Customer
Tool Product
Tool Support

Agent



LLM

Tool Customer

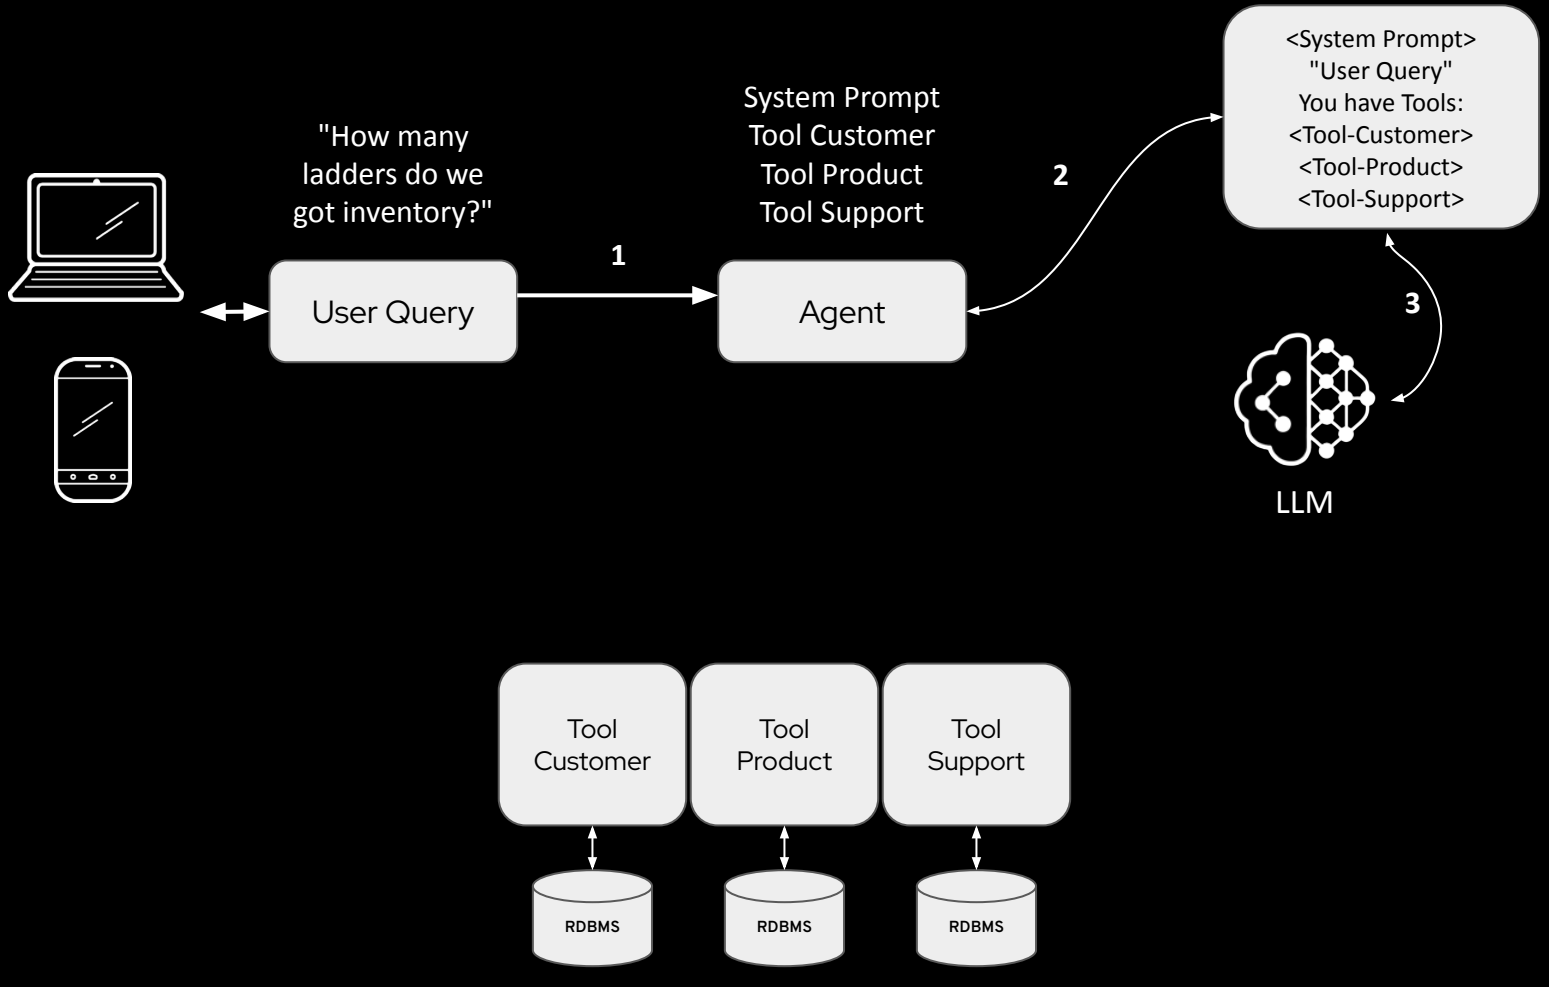
Tool Product

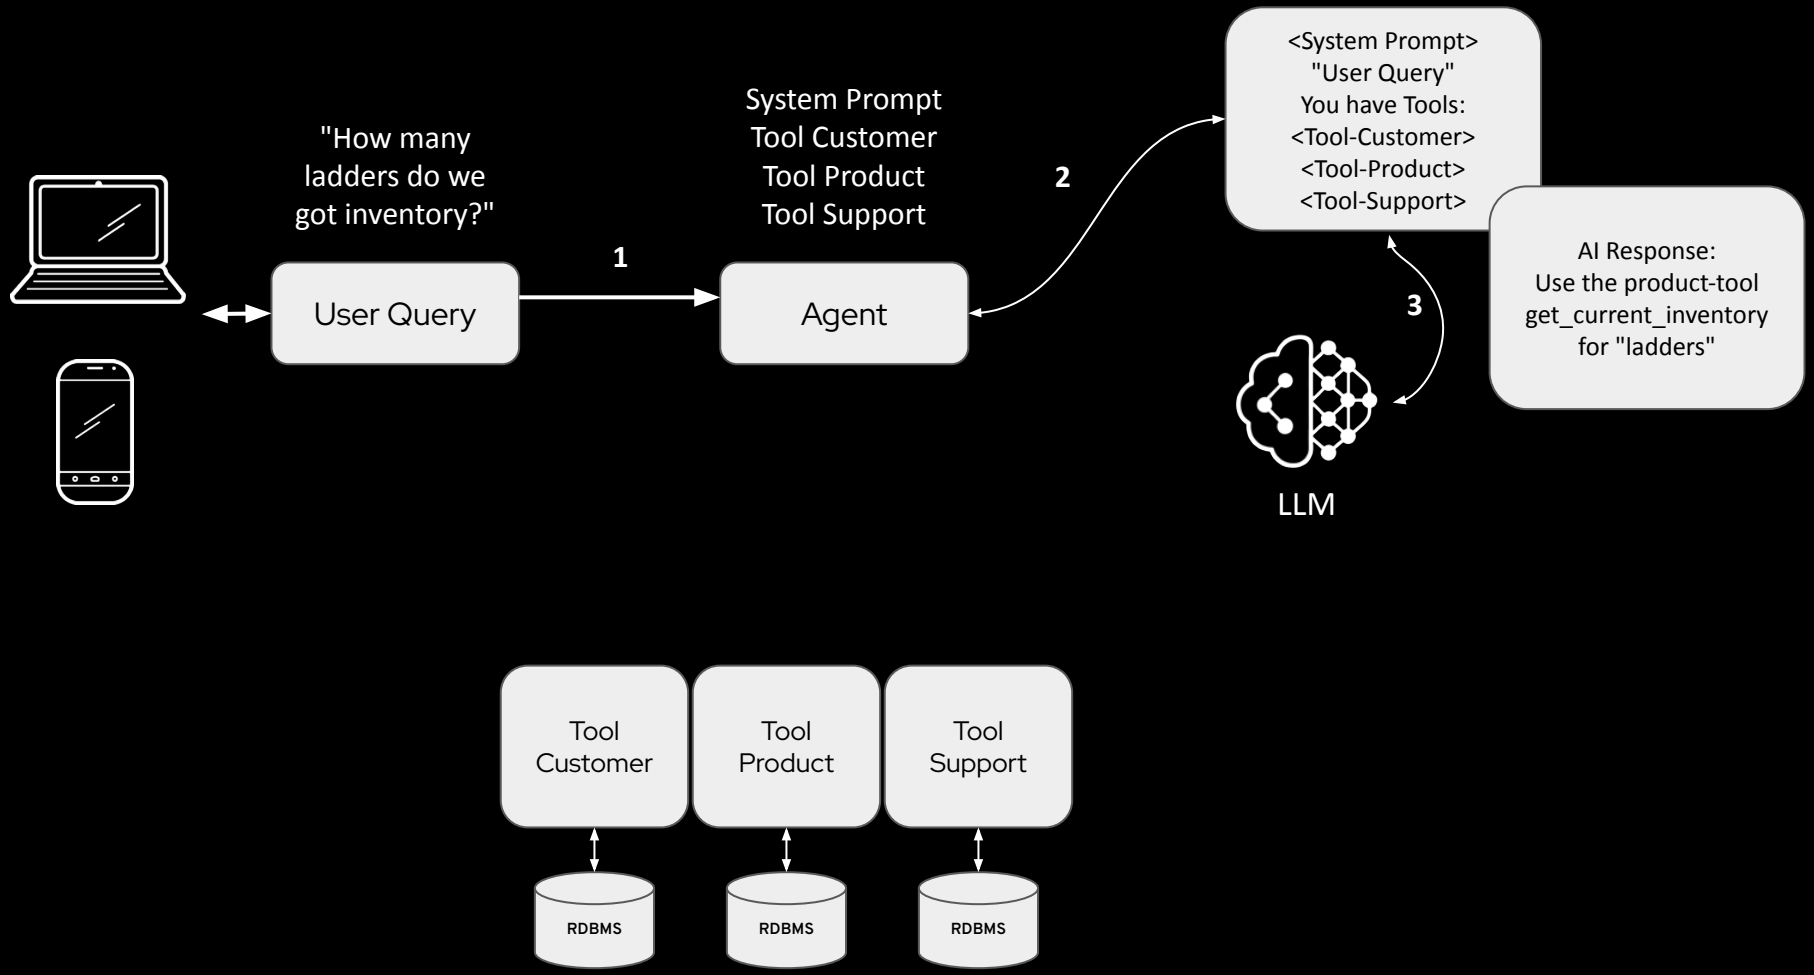
Tool Support

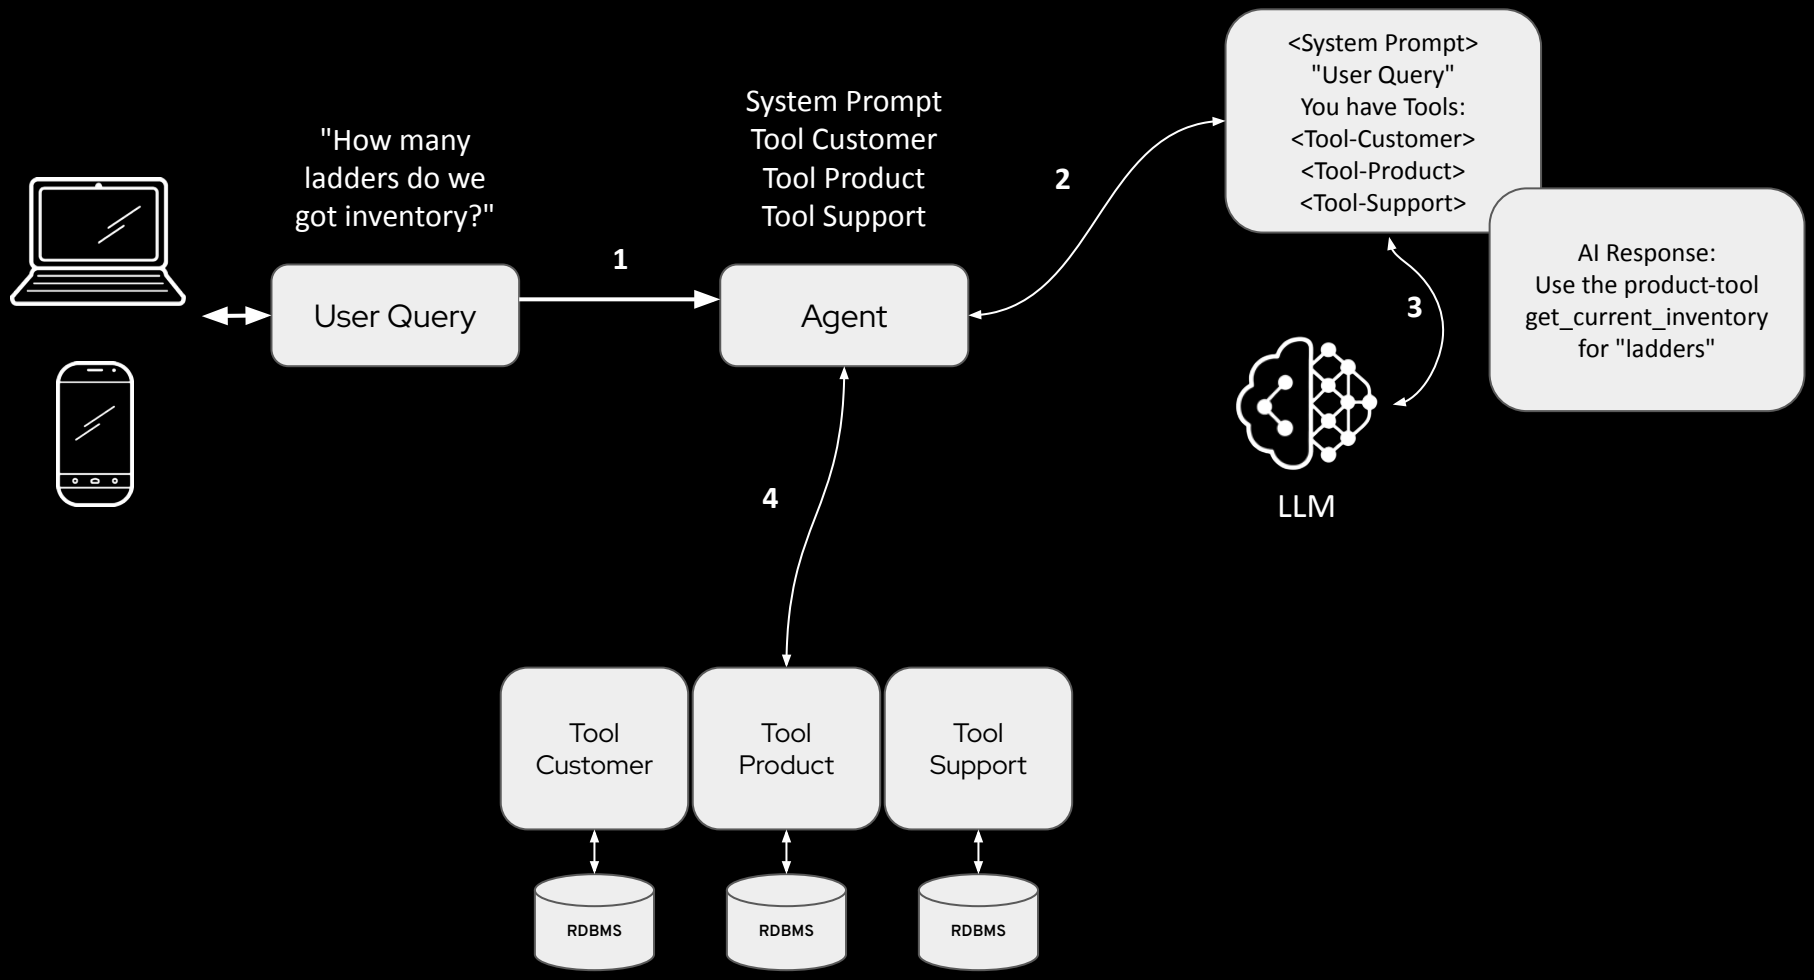
RDBMS

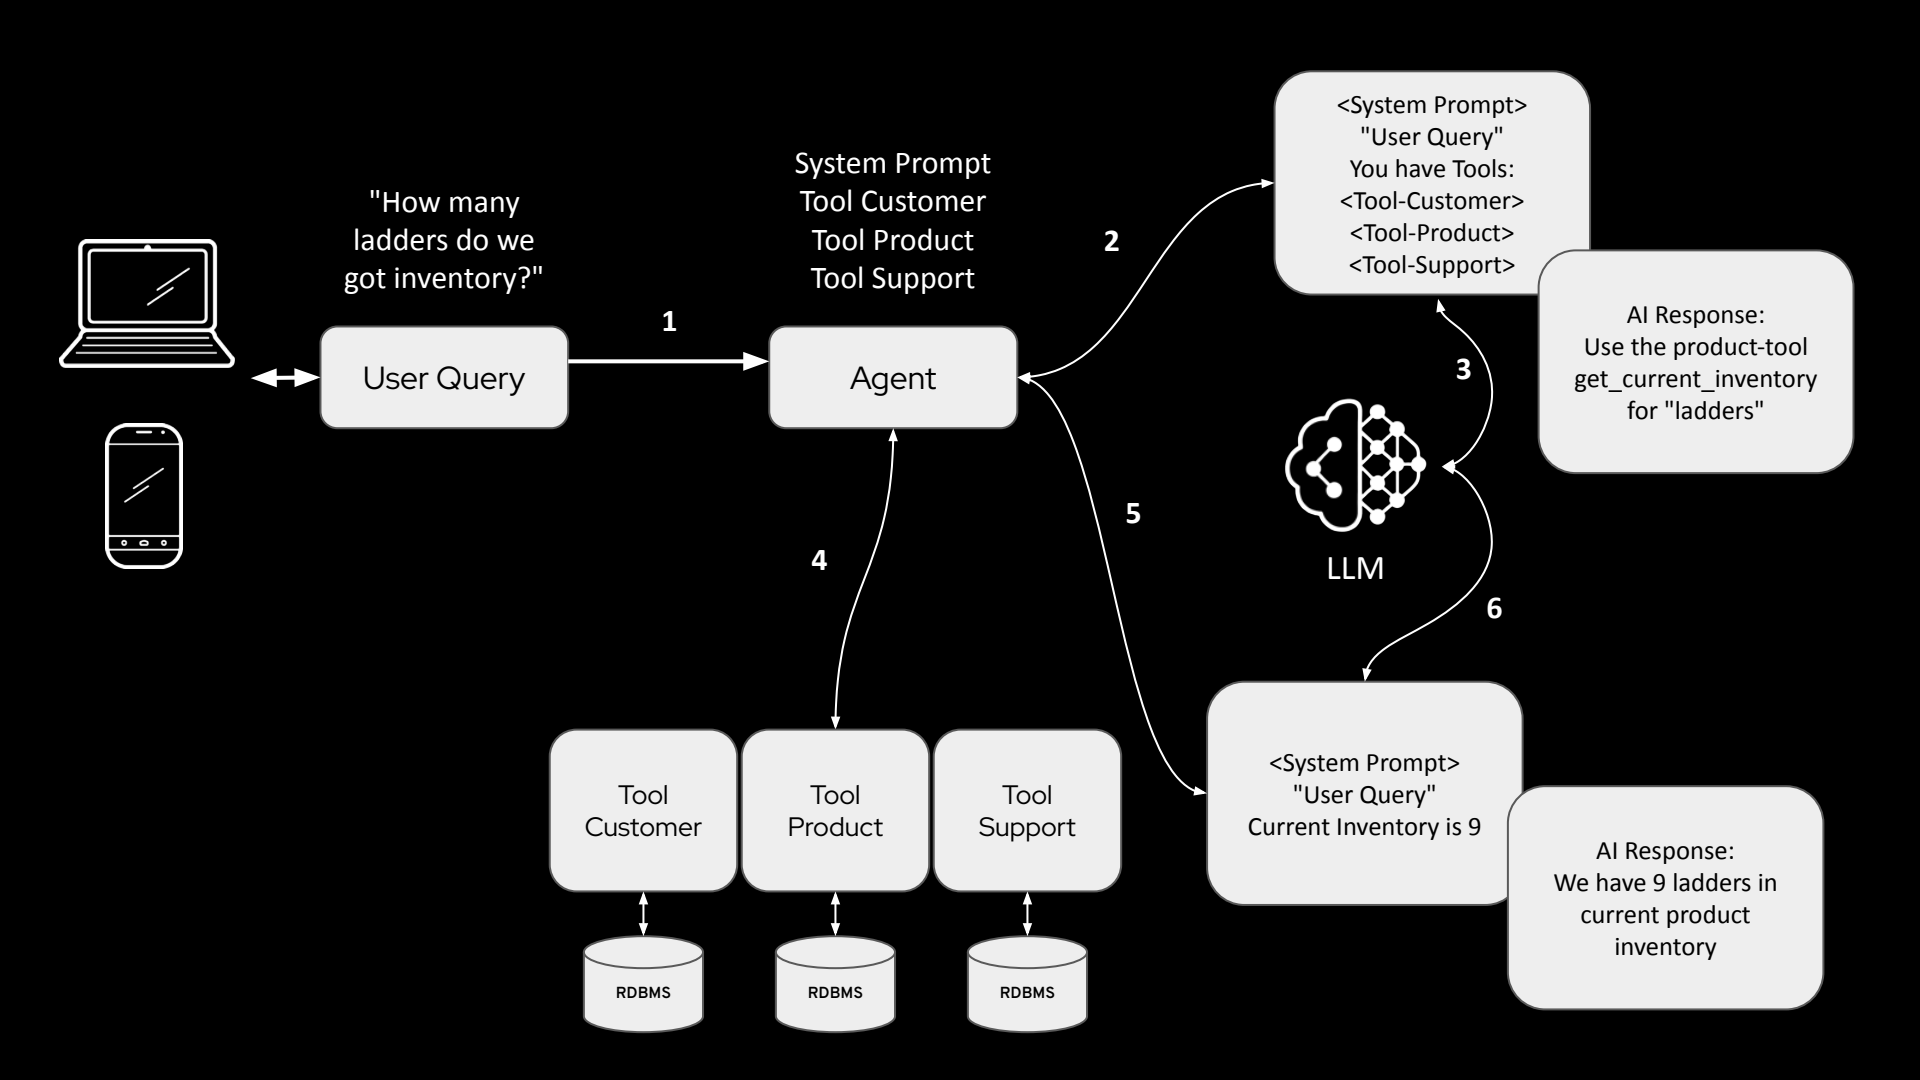
RDBMS

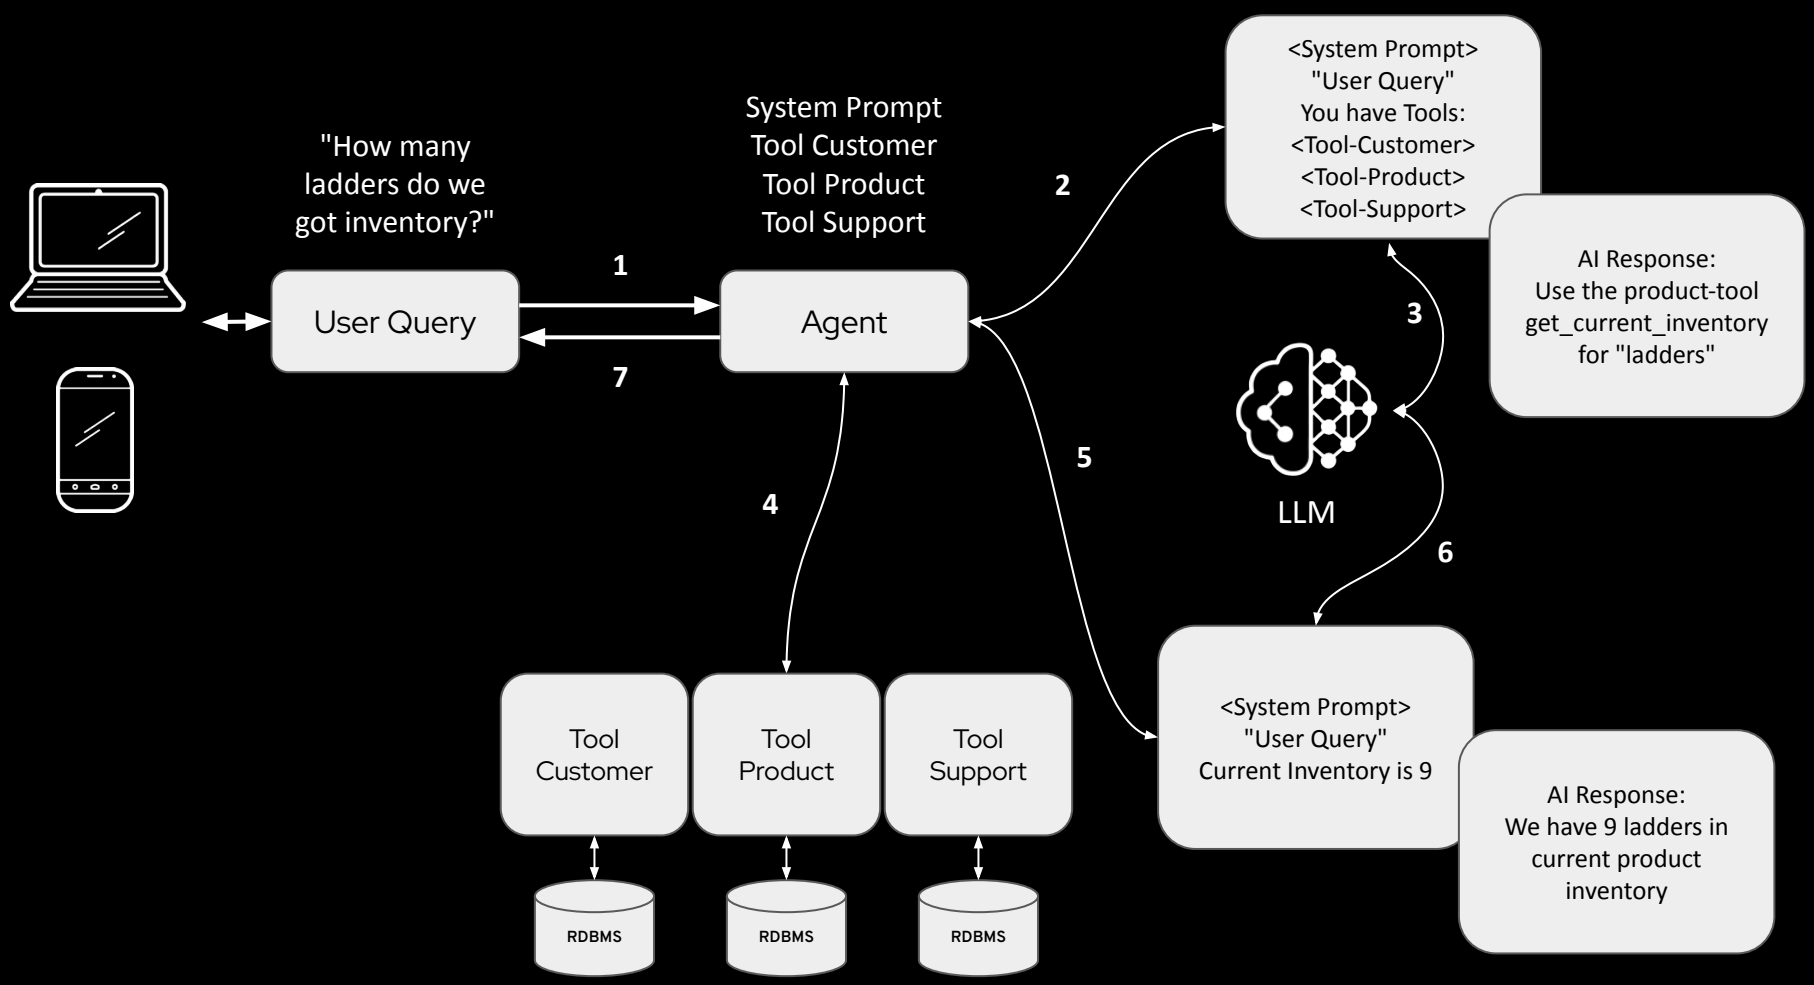
RDBMS


"How many ladders do we got inventory?"
User Query
1
System Prompt
Tool Customer
Tool Product
Tool Support
Agent
2
<System Prompt>
"User Query"
You have Tools:
<Tool-Customer>
<Tool-Product>
<Tool-Support>
3
LLM
Tool Customer
Tool Product
Tool Support
RDBMS
RDBMS
RDBMS


"How many ladders do we got inventory?"
User Query
1
System Prompt
Tool Customer
Tool Product
Tool Support
Agent
2
<System Prompt>
"User Query"
You have Tools:
<Tool-Customer>
<Tool-Product>
<Tool-Support>
3
LLM
AI Response:
Use the product-tool
get_current_inventory
for "ladders"
Tool Customer
Tool Product
Tool Support
RDBMS
RDBMS
RDBMS


"How many ladders do we got inventory?"
User Query
1
System Prompt
Tool Customer
Tool Product
Tool Support
Agent
2
<System Prompt>
"User Query"
You have Tools:
<Tool-Customer>
<Tool-Product>
<Tool-Support>
3
AI Response:
Use the product-tool
get_current_inventory
for "ladders"
LLM
4
Tool Customer
Tool Product
Tool Support
RDBMS
RDBMS
RDBMS


"How many ladders do we got inventory?"
User Query
1
System Prompt
Tool Customer
Tool Product
Tool Support
Agent
2
<System Prompt>
"User Query"
You have Tools:
<Tool-Customer>
<Tool-Product>
<Tool-Support>
3
AI Response:
Use the product-tool get_current_inventory for "ladders"
LLM
4
5
6
Tool Customer
Tool Product
Tool Support
RDBMS
RDBMS
RDBMS
<System Prompt>
"User Query"
Current Inventory is 9
AI Response:
We have 9 ladders in current product inventory


"How many ladders do we got inventory?"
User Query
1
7
System Prompt
Tool Customer
Tool Product
Tool Support
Agent
2
<System Prompt>
"User Query"
You have Tools:
<Tool-Customer>
<Tool-Product>
<Tool-Support>
3
AI Response:
Use the product-tool get_current_inventory for "ladders"
LLM
4
5
6
Tool Customer
Tool Product
Tool Support
RDBMS
RDBMS
RDBMS
<System Prompt>
"User Query"
Current Inventory is 9
AI Response:
We have 9 ladders in current product inventory

Model Context Protocol (MCP)



Tools	Resources	Prompts
Model-controlled	*Application-controlled*	*User-controlled*
Functions invoked by model	Data exposed to the application	Pre-defined templates for AI interactions
Retrieve/Search	Files	Document Q&A
Send a message	Database Records	Transcript Summary
Update DB records	API Responses	Output as JSON

Source: [Building Agents with Model Context Protocol](#) **(youtube)**

RAG
(Knowledge)

Naive RAG

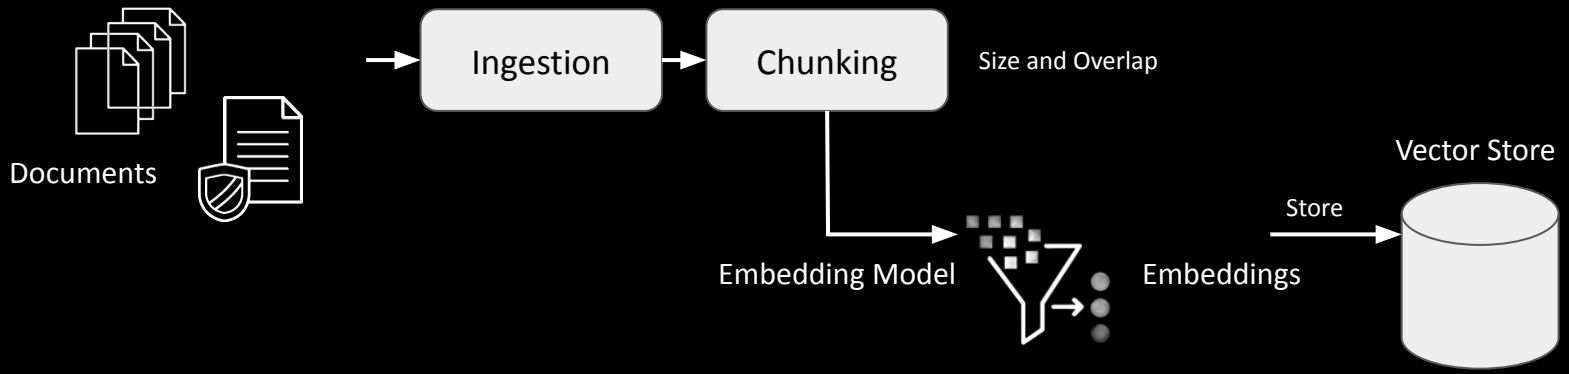


Ingestion → Chunking

Size and Overlap

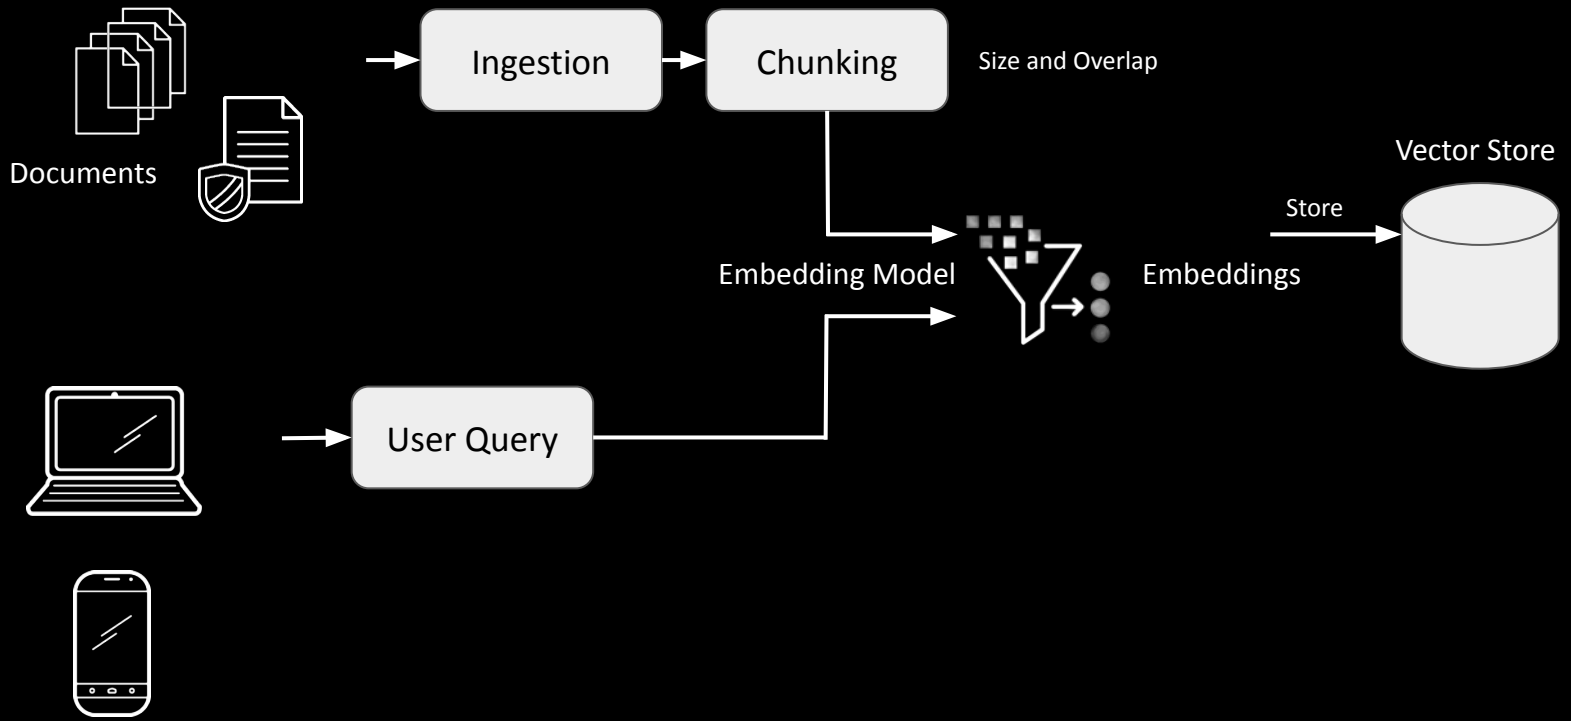
Tables, graphs or hand-writing- look at tools like Docling or VLM

Naive RAG



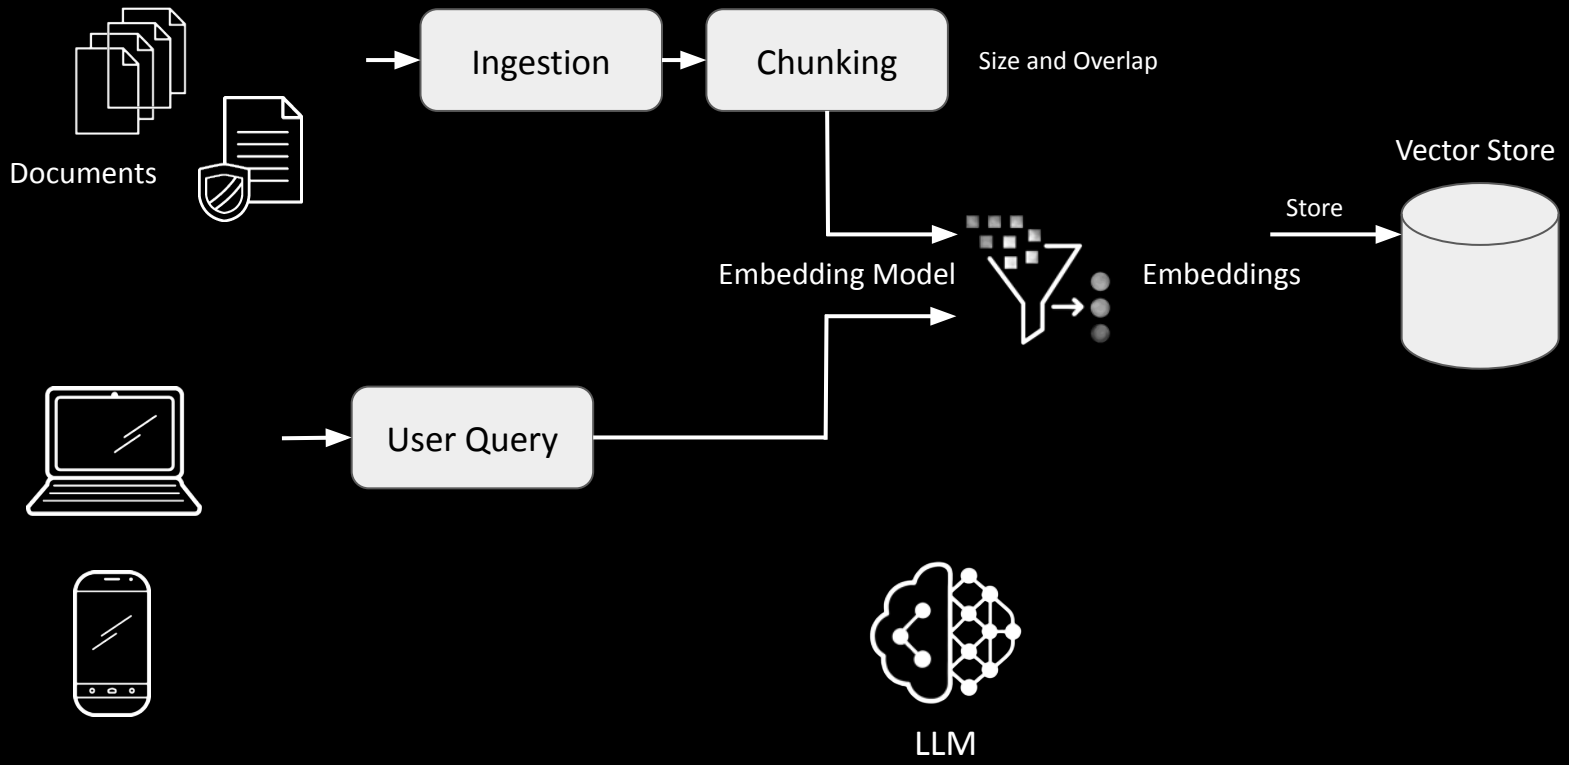
Tables, graphs or hand-writing- look at tools like Docling or VLM

Naive RAG



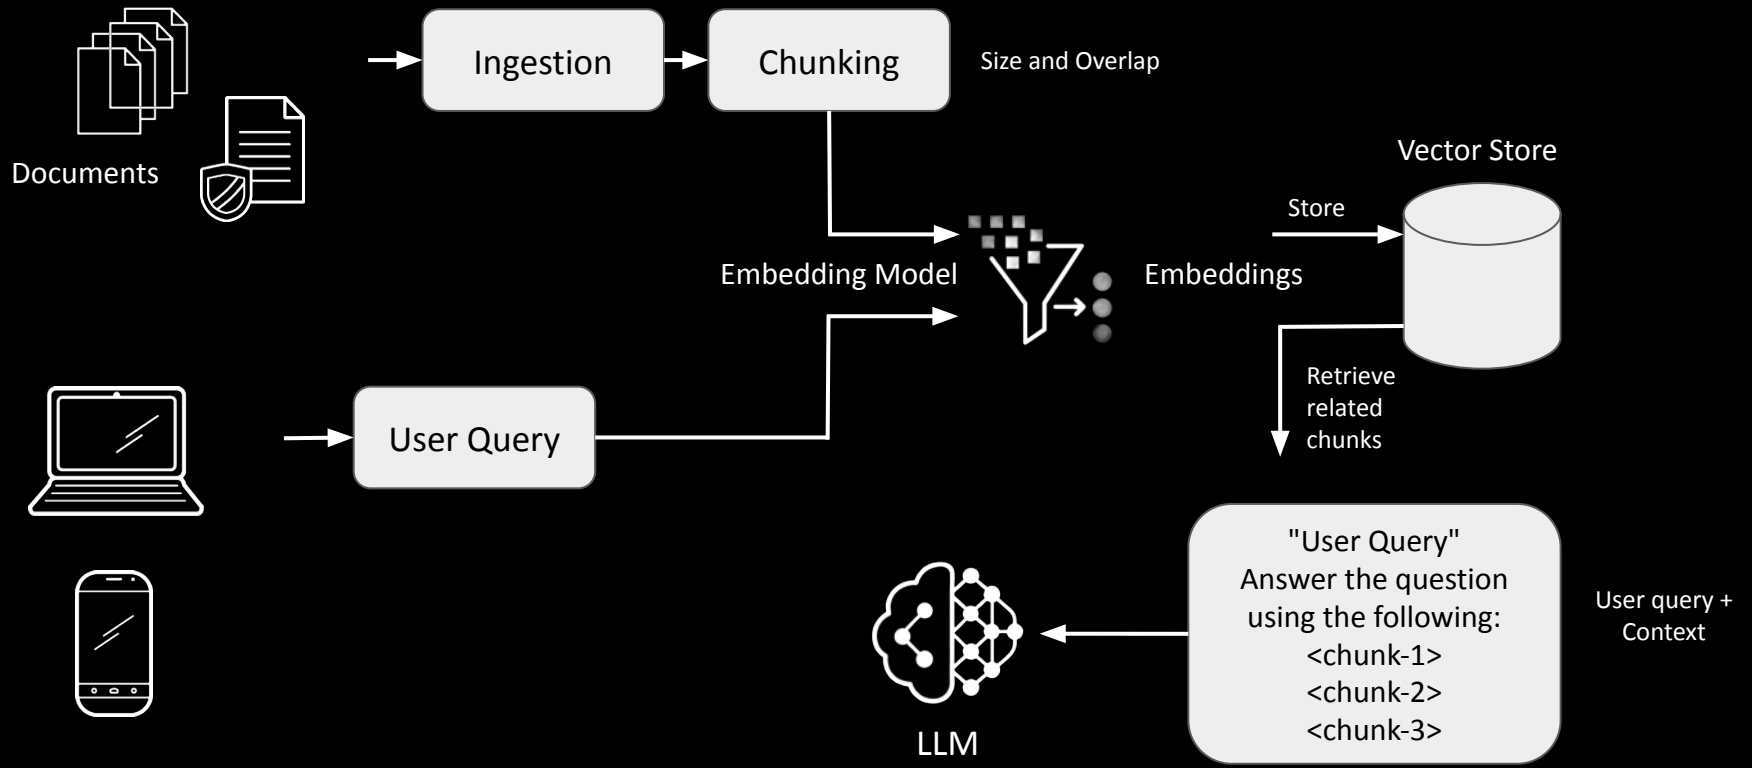
Tables, graphs or hand-writing- look at tools like Docling or VLM

Naive RAG



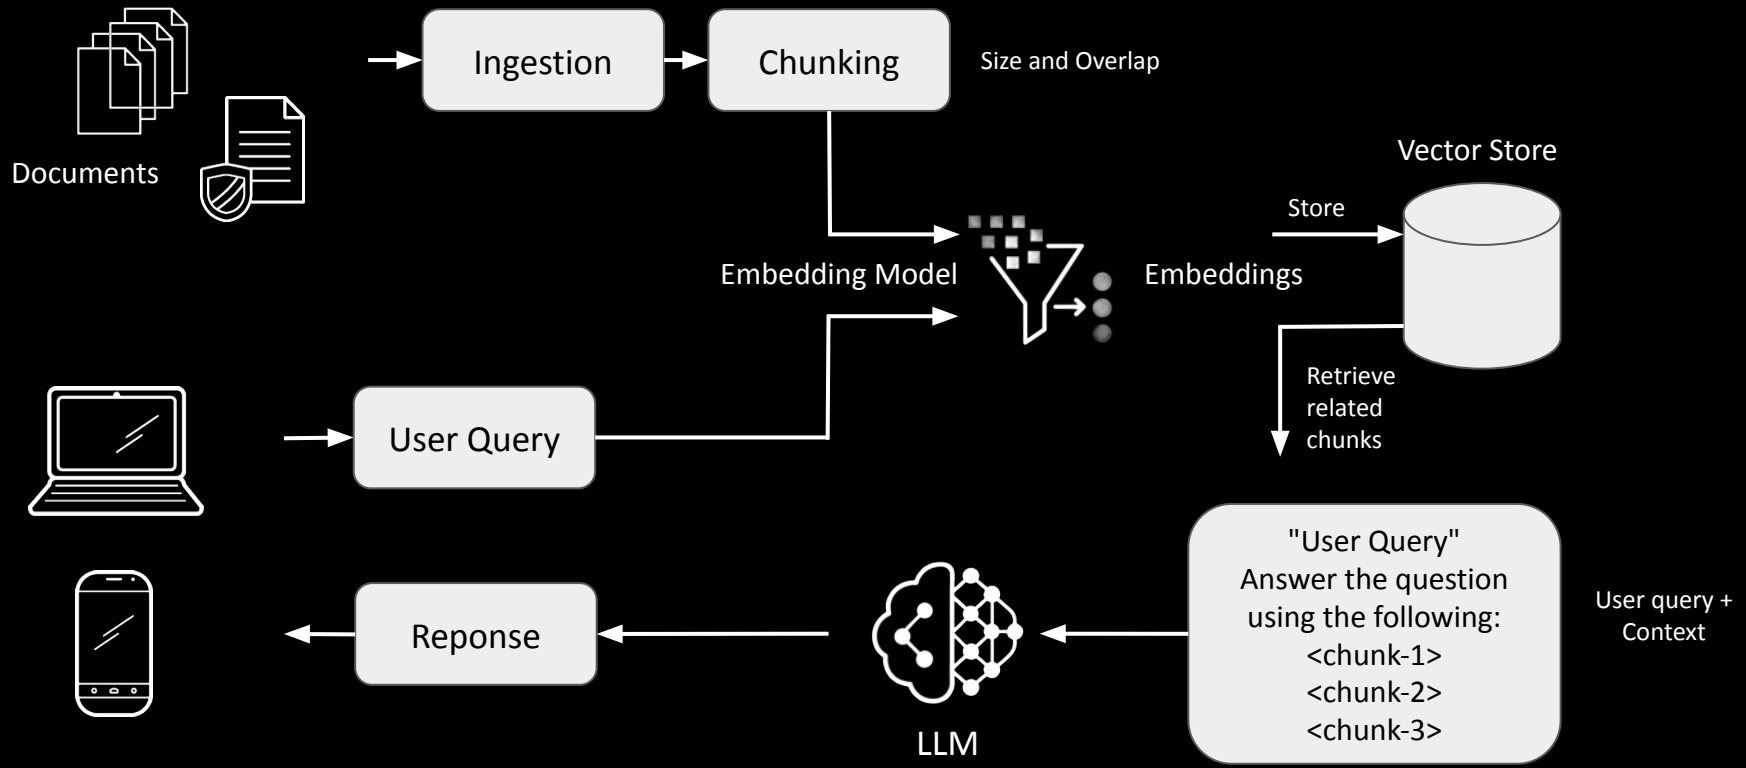
Tables, graphs or hand-writing- look at tools like Docling or VLM

Naive RAG



Tables, graphs or hand-writing- look at tools like Docling or VLM

Naive RAG



Tables, graphs or hand-writing- look at tools like Docling or VLM

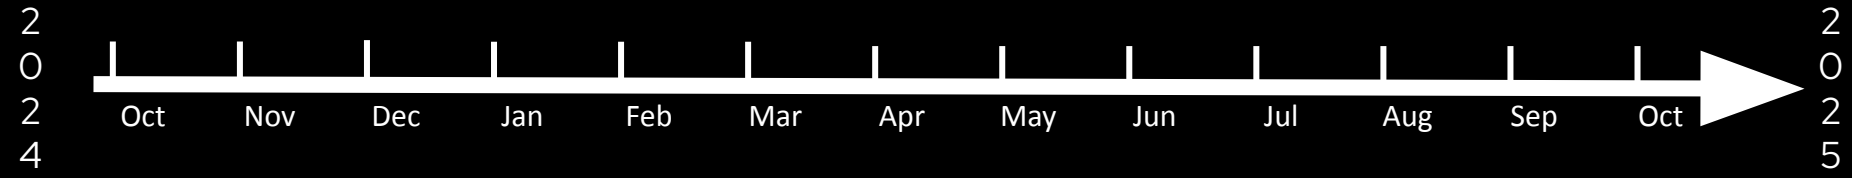
Gartner Says Worldwide AI Spending Will Total $2.5 Trillion in 2026

STAMFORD, Conn., January 15, 2026

AI Infrastructure Drives AI Spending; Adds $401 Billion in Spending as Technology Providers Continue to Build Out AI Foundations

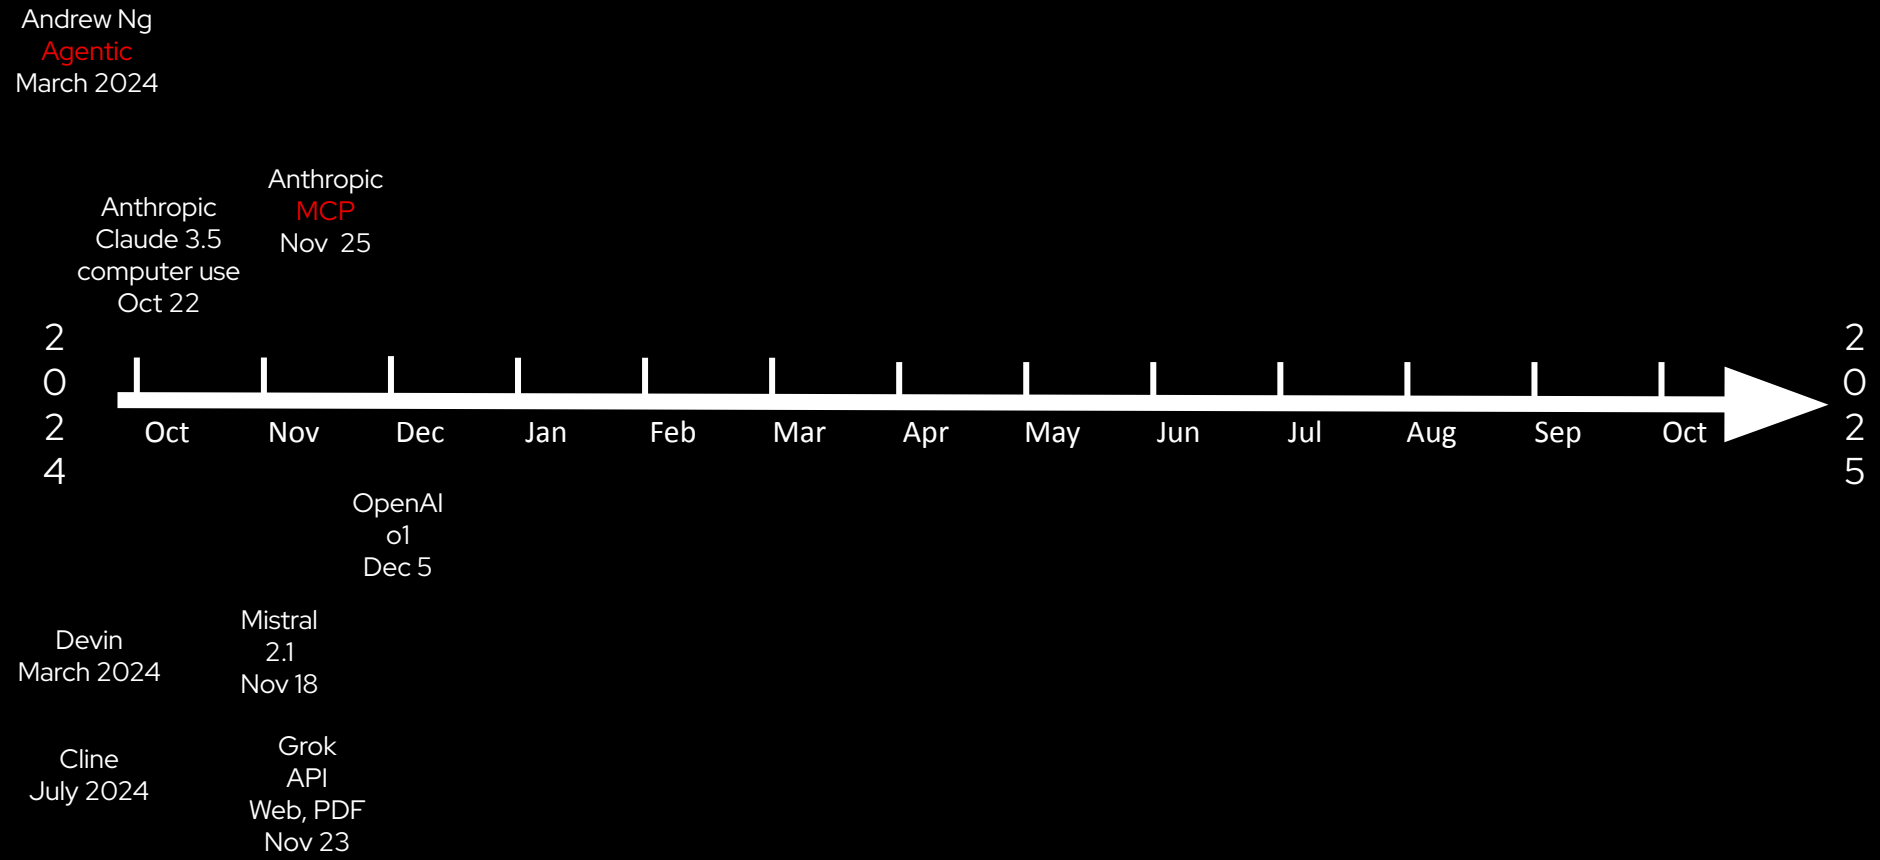
https://www.gartner.com/en/newsroom/press-releases/2026-1-15-gartner-says-worldwide-ai-spending-will-total-2-point-5-trillion-dollars-in-2026

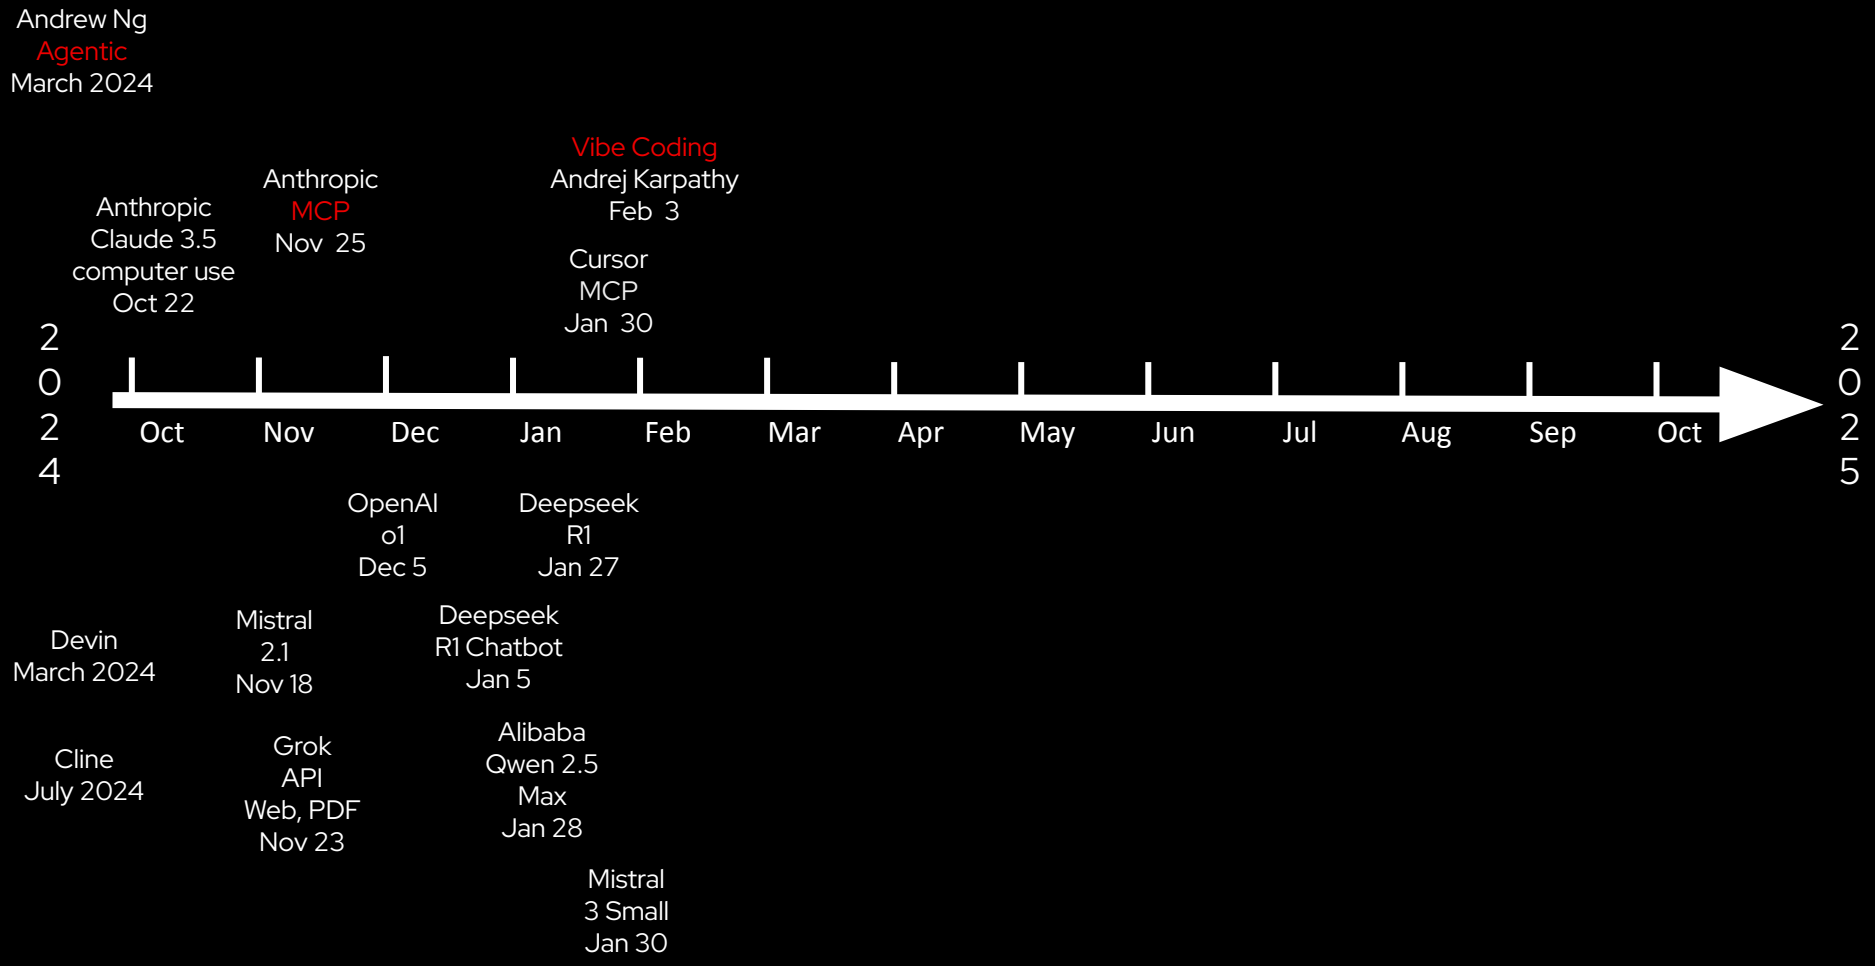
Andrew Ng
Agentic
March 2024

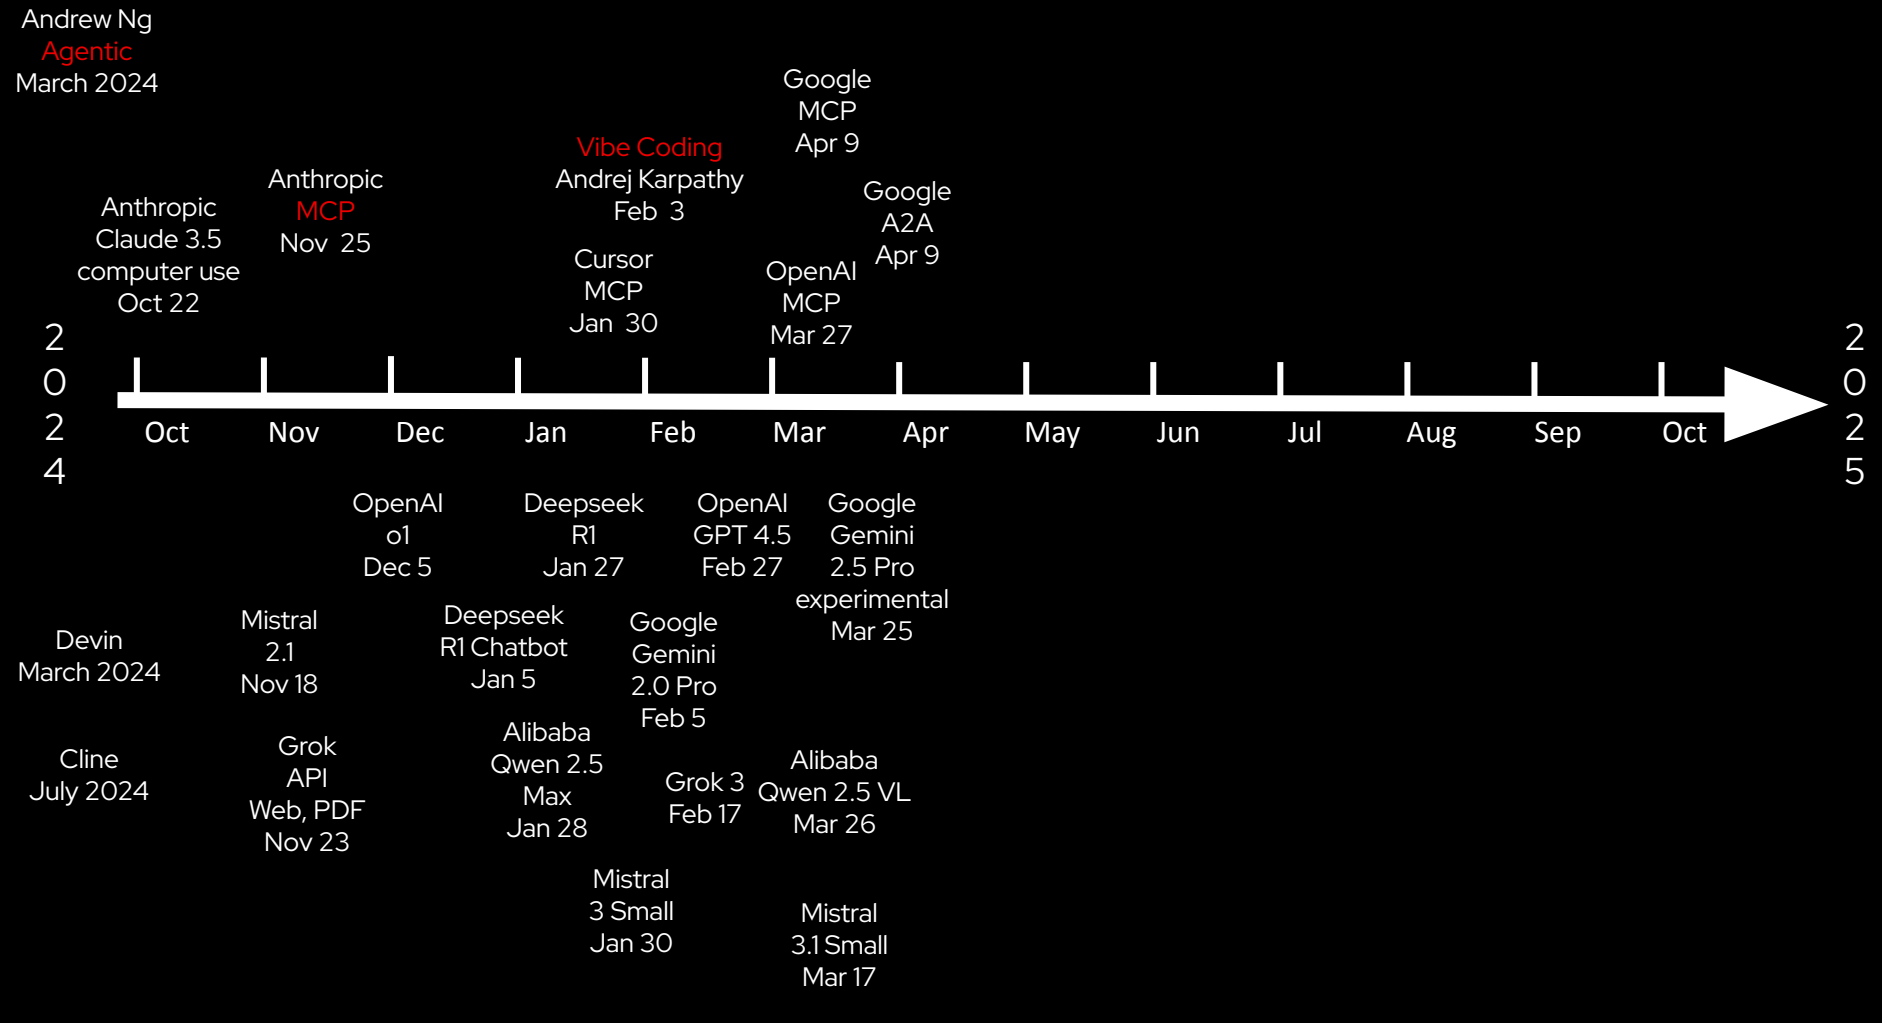


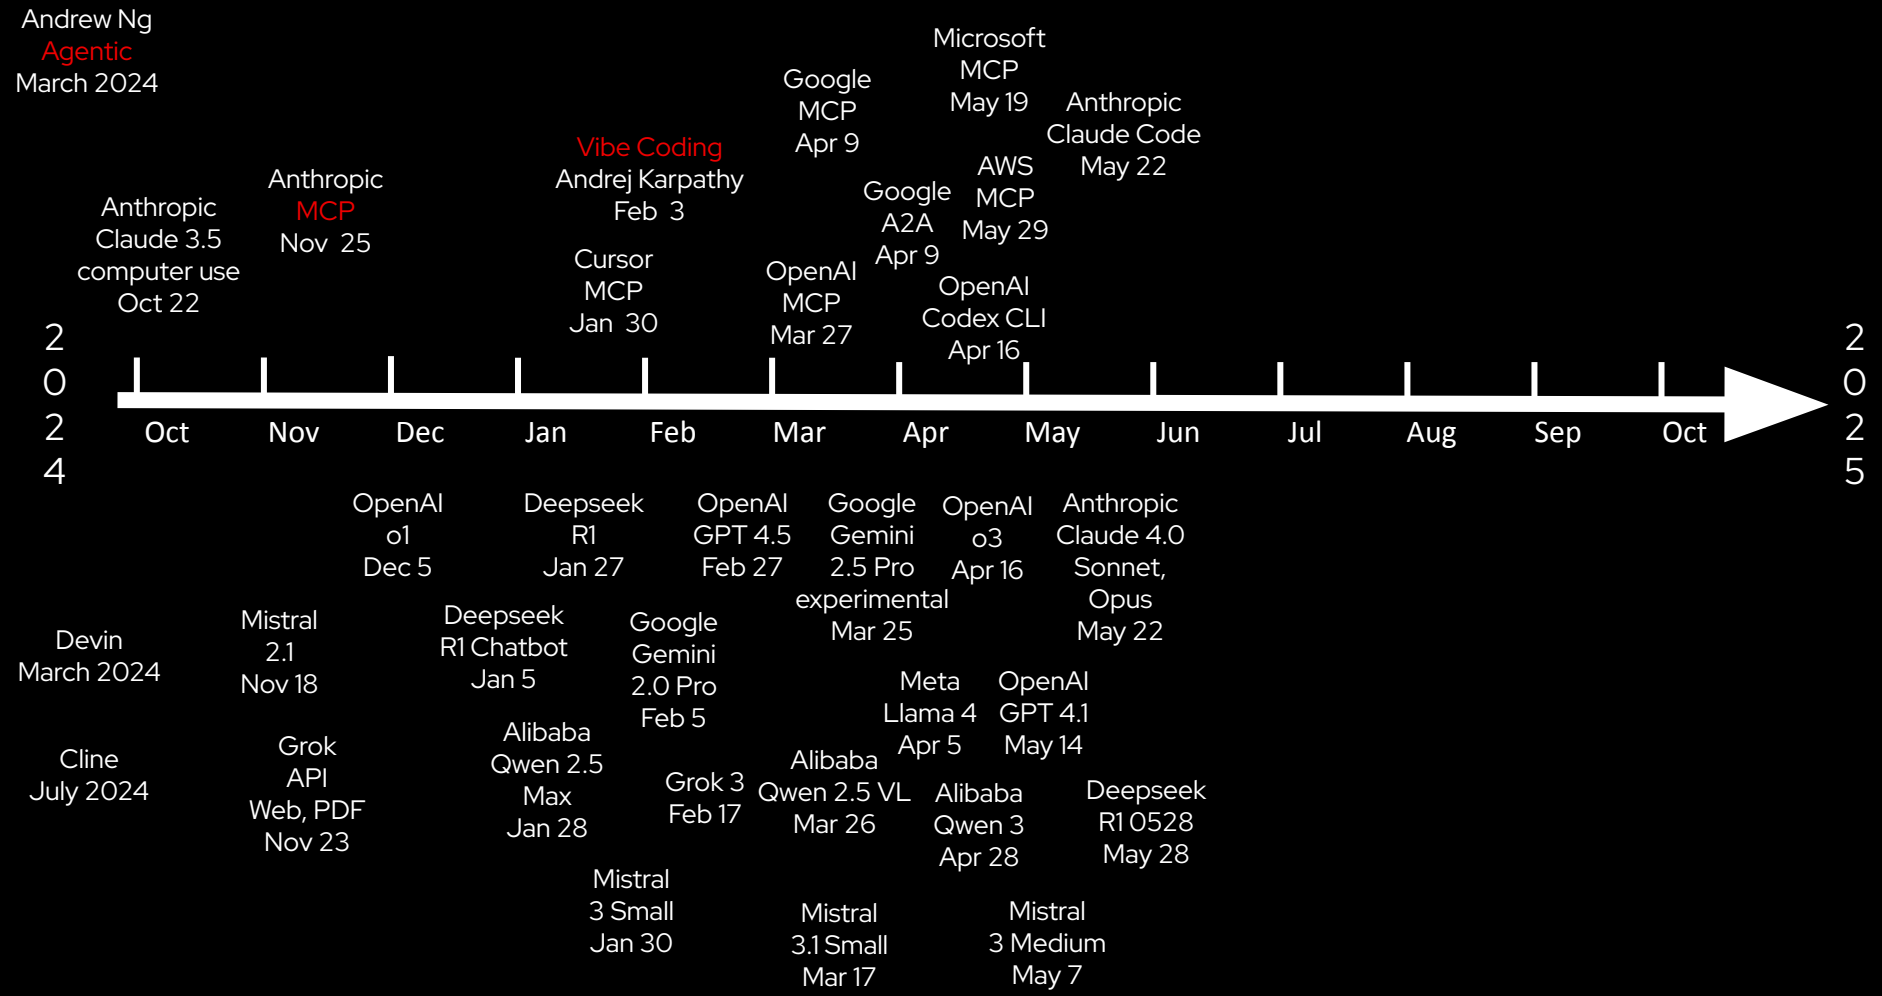
Devin
March 2024

Cline
July 2024


Andrew Ng
Agentic
March 2024
Anthropic
Claude 3.5
computer use
Oct 22
Anthropic
MCP
Nov 25
2024
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
2025
OpenAI
o1
Dec 5
Devin
March 2024
Mistral
2.1
Nov 18
Cline
July 2024
Grok
API
Web, PDF
Nov 23

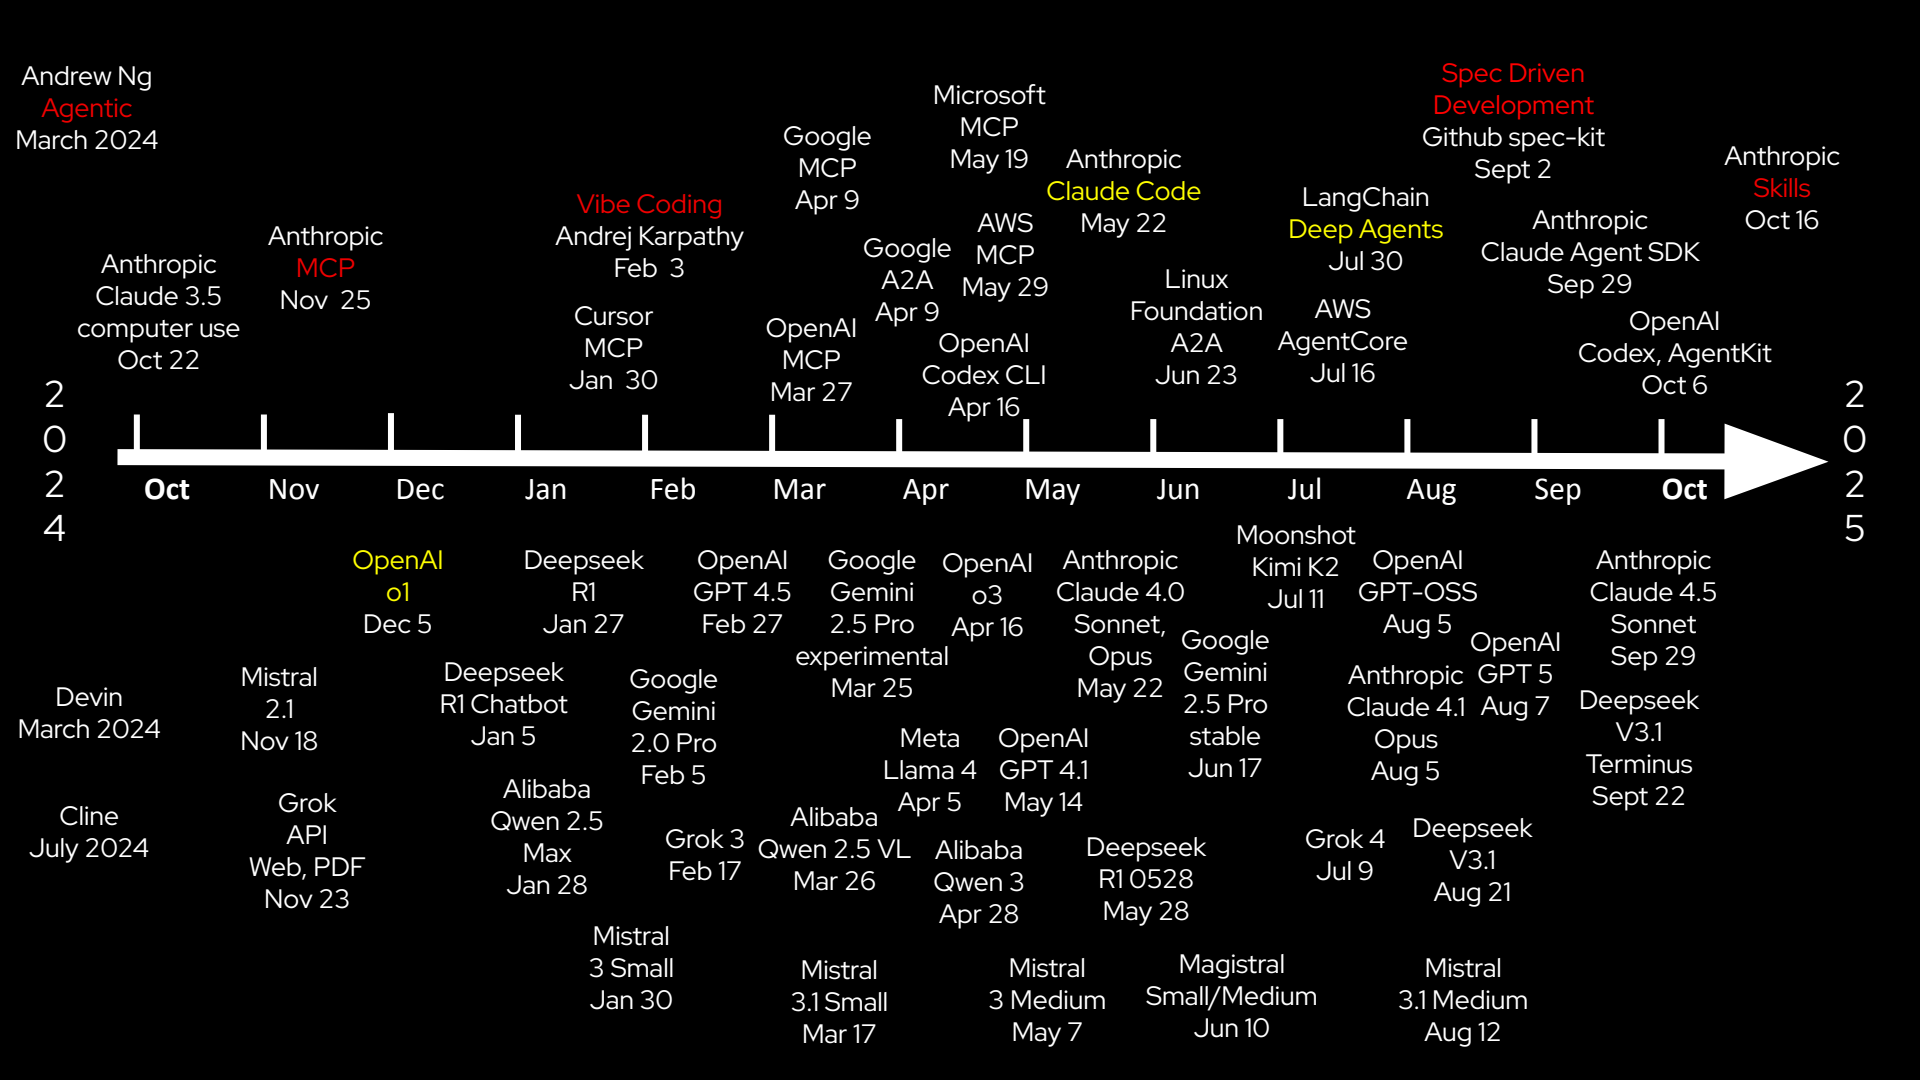

Andrew Ng
Agentic
March 2024
Anthropic
Claude 3.5
computer use
Oct 22
Anthropic
MCP
Nov 25
Vibe Coding
Andrej Karpathy
Feb 3
Cursor
MCP
Jan 30
2024
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
2025
OpenAI
o1
Dec 5
Deepseek
R1
Jan 27
Devin
March 2024
Mistral
2.1
Nov 18
Deepseek
R1 Chatbot
Jan 5
Cline
July 2024
Grok
API
Web, PDF
Nov 23
Alibaba
Qwen 2.5
Max
Jan 28
Mistral
3 Small
Jan 30

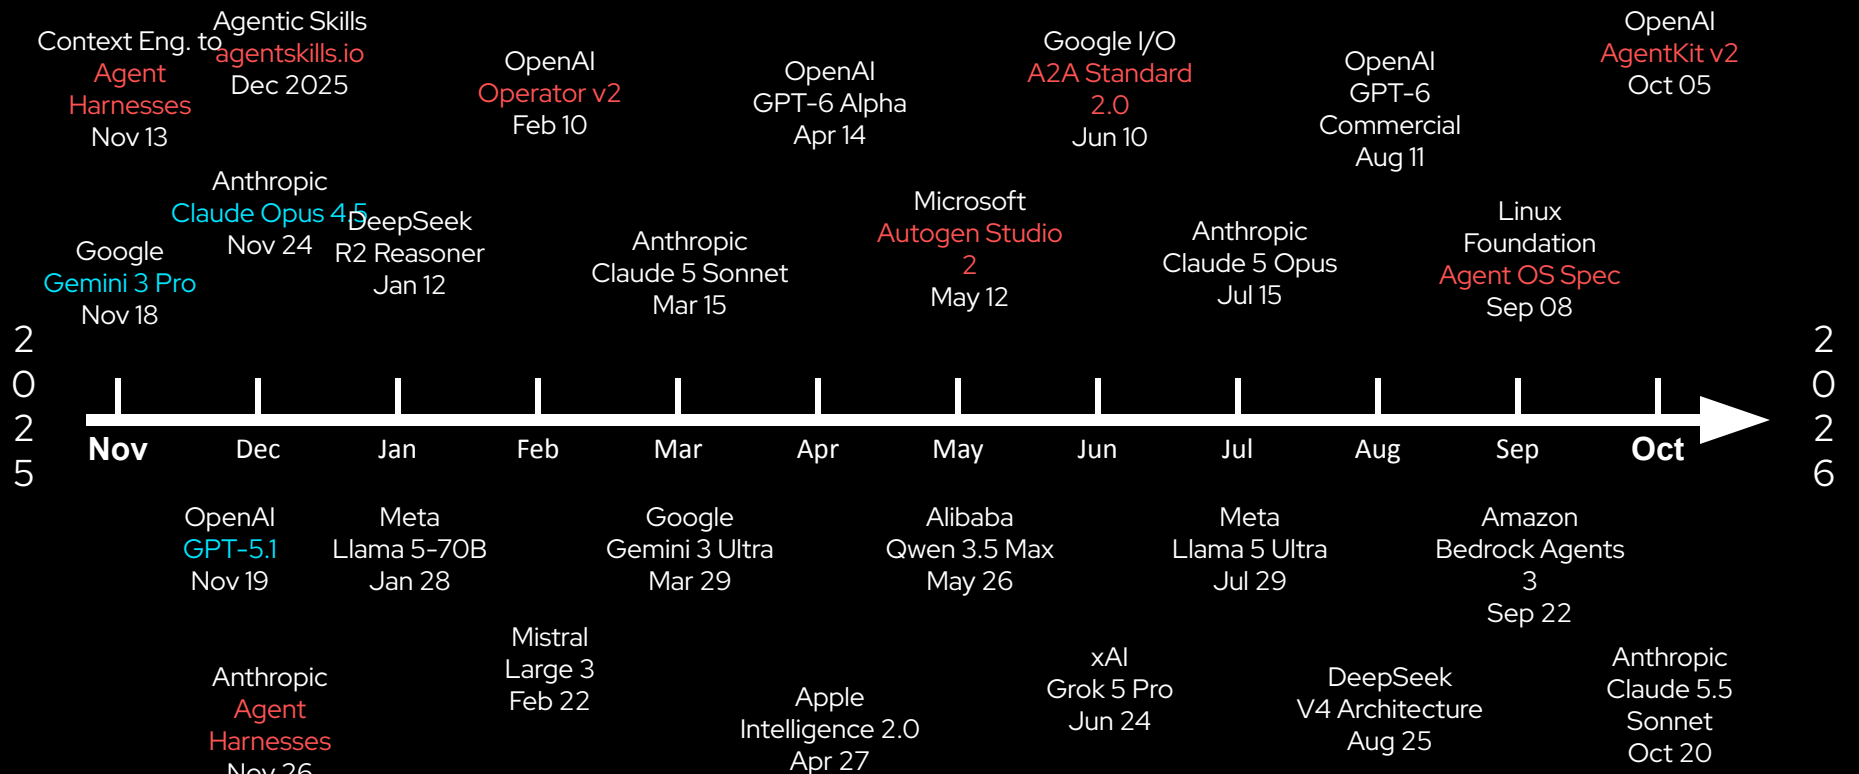

Andrew Ng
Agentic
March 2024
Anthropic
Claude 3.5
computer use
Oct 22
Anthropic
MCP
Nov 25
Vibe Coding
Andrej Karpathy
Feb 3
Cursor
MCP
Jan 30
Google
MCP
Apr 9
Google
A2A
Apr 9
OpenAI
MCP
Mar 27
2024
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
2025
OpenAI
o1
Dec 5
Deepseek
R1
Jan 27
OpenAI
GPT 4.5
Feb 27
Google
Gemini
2.5 Pro
experimental
Mar 25
Devin
March 2024
Mistral
2.1
Nov 18
Deepseek
R1 Chatbot
Jan 5
Google
Gemini
2.0 Pro
Feb 5
Cline
July 2024
Grok
API
Web, PDF
Nov 23
Alibaba
Qwen 2.5
Max
Jan 28
Grok 3
Feb 17
Alibaba
Qwen 2.5 VL
Mar 26
Mistral
3 Small
Jan 30
Mistral
3.1 Small
Mar 17


Andrew Ng
Agentic
March 2024
Anthropic
Claude 3.5
computer use
Oct 22
Anthropic
MCP
Nov 25
Vibe Coding
Andrej Karpathy
Feb 3
Cursor
MCP
Jan 30
Google
MCP
Apr 9
OpenAI
MCP
Mar 27
Google
A2A
Apr 9
Microsoft
MCP
May 19
AWS
MCP
May 29
OpenAI
Codex CLI
Apr 16
Anthropic
Claude Code
May 22
2024
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
2025
OpenAI
o1
Dec 5
Deepseek
R1
Jan 27
OpenAI
GPT 4.5
Feb 27
Google
Gemini
2.5 Pro
experimental
Mar 25
OpenAI
o3
Apr 16
Anthropic
Claude 4.0
Sonnet,
Opus
May 22
Devin
March 2024
Mistral
2.1
Nov 18
Deepseek
R1 Chatbot
Jan 5
Google
Gemini
2.0 Pro
Feb 5
Meta
Llama 4
Apr 5
OpenAI
GPT 4.1
May 14
Cline
July 2024
Grok
API
Web, PDF
Nov 23
Alibaba
Qwen 2.5
Max
Jan 28
Grok 3
Feb 17
Alibaba
Qwen 2.5 VL
Mar 26
Alibaba
Qwen 3
Apr 28
Deepseek
R1 0528
May 28
Mistral
3 Small
Jan 30
Mistral
3.1 Small
Mar 17
Mistral
3 Medium
May 7


Andrew Ng
Agentic
March 2024
Anthropic
Claude 3.5
computer use
Oct 22
Anthropic
MCP
Nov 25
Vibe Coding
Andrej Karpathy
Feb 3
Cursor
MCP
Jan 30
Google
MCP
Apr 9
OpenAI
MCP
Mar 27
Microsoft
MCP
May 19
Google
A2A
Apr 9
AWS
MCP
May 29
OpenAI
Codex CLI
Apr 16
Anthropic
Claude Code
May 22
Linux
Foundation
A2A
Jun 23
AWS
AgentCore
Jul 16
2024
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
2025
OpenAI
o1
Dec 5
Deepseek
R1
Jan 27
OpenAI
GPT 4.5
Feb 27
Google
Gemini
2.5 Pro
experimental
Mar 25
OpenAI
o3
Apr 16
Anthropic
Claude 4.0
Sonnet,
Opus
May 22
Moonshot
Kimi K2
Jul 11
OpenAI
GPT-OSS
Aug 5
Devin
March 2024
Mistral
2.1
Nov 18
Deepseek
R1 Chatbot
Jan 5
Google
Gemini
2.0 Pro
Feb 5
Meta
Llama 4
Apr 5
OpenAI
GPT 4.1
May 14
Google
Gemini
2.5 Pro
stable
Jun 17
Anthropic
Claude 4.1
Opus
Aug 5
Cline
July 2024
Grok
API
Web, PDF
Nov 23
Alibaba
Qwen 2.5
Max
Jan 28
Grok 3
Feb 17
Alibaba
Qwen 2.5 VL
Mar 26
Alibaba
Qwen 3
Apr 28
Deepseek
R1 0528
May 28
Grok 4
Jul 9
Mistral
3 Small
Jan 30
Mistral
3.1 Small
Mar 17
Mistral
3 Medium
May 7
Magistral
Small/Medium
Jun 10


Andrew Ng
Agentic
March 2024
Anthropic
Claude 3.5
computer use
Oct 22
Anthropic
MCP
Nov 25
Vibe Coding
Andrej Karpathy
Feb 3
Cursor
MCP
Jan 30
Google
MCP
Apr 9
OpenAI
MCP
Mar 27
Microsoft
MCP
May 19
Google
A2A
Apr 9
AWS
MCP
May 29
OpenAI
Codex CLI
Apr 16
Anthropic
Claude Code
May 22
Linux
Foundation
A2A
Jun 23
LangChain
Deep Agents
Jul 30
AWS
AgentCore
Jul 16
Spec Driven
Development
Github spec-kit
Sept 2
Anthropic
Claude Agent SDK
Sep 29
OpenAI
Codex, AgentKit
Oct 6
Anthropic
Skills
Oct 16
2024
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
2025
OpenAI
o1
Dec 5
Deepseek
R1
Jan 27
OpenAI
GPT 4.5
Feb 27
Google
Gemini
2.5 Pro
experimental
Mar 25
OpenAI
o3
Apr 16
Anthropic
Claude 4.0
Sonnet,
Opus
May 22
Moonshot
Kimi K2
Jul 11
OpenAI
GPT-OSS
Aug 5
Anthropic
Claude 4.5
Sonnet
Sep 29
Devin
March 2024
Mistral
2.1
Nov 18
Deepseek
R1 Chatbot
Jan 5
Google
Gemini
2.0 Pro
Feb 5
Meta
Llama 4
Apr 5
OpenAI
GPT 4.1
May 14
Google
Gemini
2.5 Pro
stable
Jun 17
Anthropic
Claude 4.1
Opus
Aug 5
OpenAI
GPT 5
Aug 7
Deepseek
V3.1
Terminus
Sept 22
Cline
July 2024
Grok
API
Web, PDF
Nov 23
Alibaba
Qwen 2.5
Max
Jan 28
Grok 3
Feb 17
Alibaba
Qwen 2.5 VL
Mar 26
Alibaba
Qwen 3
Apr 28
Deepseek
R1 0528
May 28
Grok 4
Jul 9
Deepseek
V3.1
Aug 21
Mistral
3 Small
Jan 30
Mistral
3.1 Small
Mar 17
Mistral
3 Medium
May 7
Magistral
Small/Medium
Jun 10
Mistral
3.1 Medium
Aug 12


Andrew Ng
Agentic
March 2024
Anthropic
Claude 3.5
computer use
Oct 22
Anthropic
MCP
Nov 25
Vibe Coding
Andrej Karpathy
Feb 3
Cursor
MCP
Jan 30
Google
MCP
Apr 9
OpenAI
MCP
Mar 27
Microsoft
MCP
May 19
Google
A2A
Apr 9
AWS
MCP
May 29
OpenAI
Codex CLI
Apr 16
Anthropic
Claude Code
May 22
Linux
Foundation
A2A
Jun 23
LangChain
Deep Agents
Jul 30
AWS
AgentCore
Jul 16
Spec Driven
Development
Github spec-kit
Sept 2
Anthropic
Claude Agent SDK
Sep 29
OpenAI
Codex, AgentKit
Oct 6
Anthropic
Skills
Oct 16
2024
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
2025
OpenAI
o1
Dec 5
Deepseek
R1
Jan 27
OpenAI
GPT 4.5
Feb 27
Google
Gemini
2.5 Pro
experimental
Mar 25
OpenAI
o3
Apr 16
Anthropic
Claude 4.0
Sonnet,
Opus
May 22
Moonshot
Kimi K2
Jul 11
OpenAI
GPT-OSS
Aug 5
Anthropic
Claude 4.5
Sonnet
Sep 29
Devin
March 2024
Mistral
2.1
Nov 18
Deepseek
R1 Chatbot
Jan 5
Google
Gemini
2.0 Pro
Feb 5
Meta
Llama 4
Apr 5
OpenAI
GPT 4.1
May 14
Google
Gemini
2.5 Pro
stable
Jun 17
Anthropic
Claude 4.1
Opus
Aug 5
OpenAI
GPT 5
Aug 7
Deepseek
V3.1
Terminus
Sept 22
Cline
July 2024
Grok
API
Web, PDF
Nov 23
Alibaba
Qwen 2.5
Max
Jan 28
Grok 3
Feb 17
Alibaba
Qwen 2.5 VL
Mar 26
Alibaba
Qwen 3
Apr 28
Deepseek
R1 0528
May 28
Grok 4
Jul 9
Deepseek
V3.1
Aug 21
Mistral
3 Small
Jan 30
Mistral
3.1 Small
Mar 17
Mistral
3 Medium
May 7
Magistral
Small/Medium
Jun 10
Mistral
3.1 Medium
Aug 12


Context Eng. to Agent Harnesses Nov 13
Agentic Skills agentskills.io Dec 2025
OpenAI Operator v2 Feb 10
OpenAI GPT-6 Alpha Apr 14
Google I/O A2A Standard 2.0 Jun 10
OpenAI GPT-6 Commercial Aug 11
OpenAI AgentKit v2 Oct 05
Google Gemini 3 Pro Nov 18
Anthropic Claude Opus 4.5 Nov 24
DeepSeek R2 Reasoner Jan 12
Anthropic Claude 5 Sonnet Mar 15
Microsoft Autogen Studio 2 May 12
Anthropic Claude 5 Opus Jul 15
Linux Foundation Agent OS Spec Sep 08
2025
2026
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
OpenAI GPT-5.1 Nov 19
Meta Llama 5-70B Jan 28
Google Gemini 3 Ultra Mar 29
Alibaba Qwen 3.5 Max May 26
Meta Llama 5 Ultra Jul 29
Amazon Bedrock Agents 3 Sep 22
Anthropic Agent Harnesses Nov 26
Mistral Large 3 Feb 22
Apple Intelligence 2.0 Apr 27
xAI Grok 5 Pro Jun 24
DeepSeek V4 Architecture Aug 25
Anthropic Claude 5.5 Sonnet Oct 20

agentskills.io

SKILL.md
evals.json



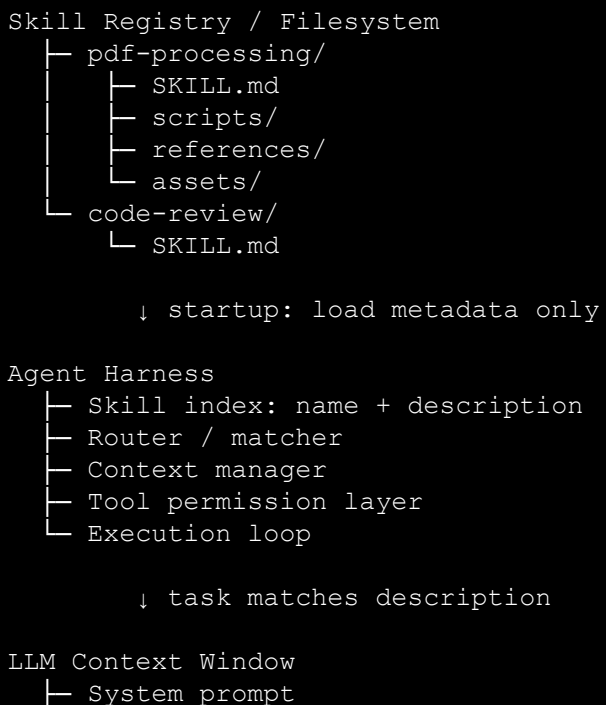




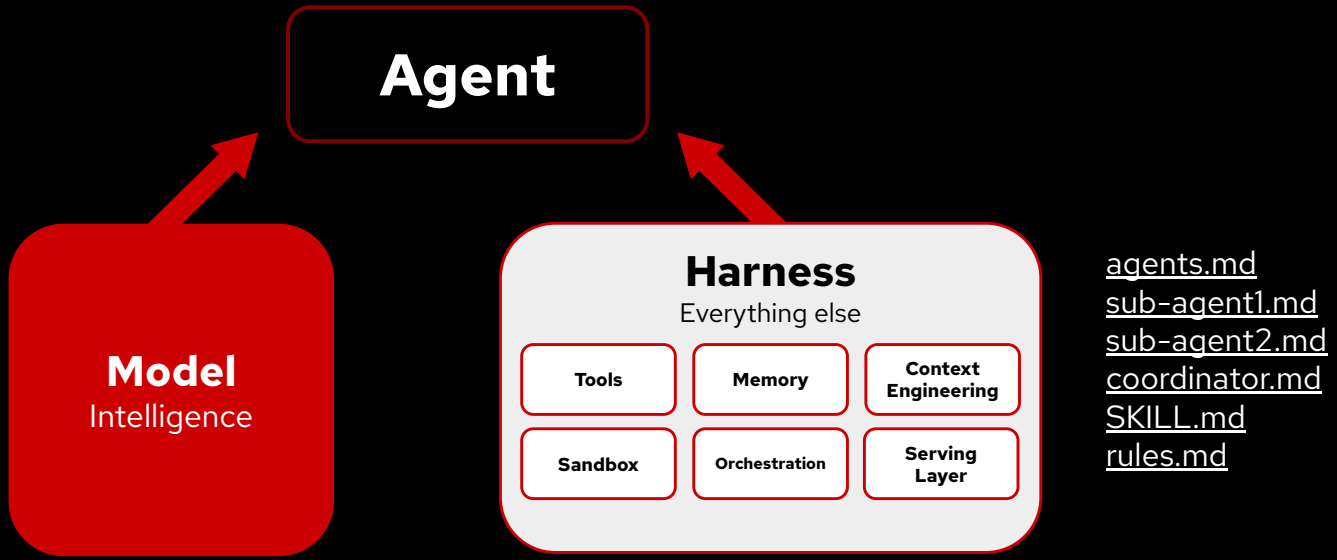

SKILL.md

Skills (agentskills.io)

- Progressive Discovery
  - Unlike Tools/MCP, skills are loaded dynamically as needed
  - Only name/description load at startup
  - Address system prompt bloat (and saves you tokens)
- Cross Platform Compatibility
  - Works across Claude Code, Codex, ChatGPT, Cursor, GitHub Copilot, OpenClaw, and 26+ platforms
  - Open standard format & part of vendor neutral Agentic AI Foundation (AAIF)
- Markdown Simplicity
- Executable Code
  - Bundle scripts in the `scripts/` sub-folder
- On-Demand Domain Expertise
  - Packages procedural knowledge and company/team/user-specific context


Skill Registry / Filesystem
├─ pdf-processing/
│ ├─ SKILL.md
│ ├─ scripts/
│ ├─ references/
│ └─ assets/
└─ code-review/
└─ SKILL.md
↓ startup: load metadata only
Agent Harness
├─ Skill index: name + description
├─ Router / matcher
├─ Context manager
├─ Tool permission layer
└─ Execution loop
↓ task matches description
LLM Context Window
├─ System prompt
├─ User request
├─ Selected SKILL.md
└─ Optional referenced files/scripts output

The Rise of **Harness Engineering**



agents.md
sub-agent1.md
sub-agent2.md
coordinator.md
SKILL.md
rules.md

Agent Harness “Factors”

1. Agentic Loop
2. Context Management
3. Tools
4. Skills
5. Subagent Delegation
6. Memory (Working + Persistent)
7. Permissions & Security & Guardrails
8. Observability & Evals
9. Error Recovery & Resilience
10. System Prompt & Instruction Hierarchy
11. Lifecycle Hooks



Claude Code, Codex, Deep Agents,
Cursor, Antigravity,
OpenClaw, OpenCode

Deep Agents

from LangChain

Cursor
Claude Code
OpenClaw
Codex
Open Code

Enterprise

- Systems level automations (Potentially dozens of Agents)
- Hundreds of users per Agent | Many Skills per Agent

Domain / Org / Department

- Multi-human automations (Potentially hundreds of Agents)

Team

- Multi-human automations (Potentially thousands of Agents)

Personal

- 20K Employees → 20K Agents (At least one per Human)
- 100K+ Personal Skills | Multiple Skills per Agent

Tribal Knowledge written down as `Skill.md`

The Agent Harness = Operating System

Same architecture, different era

COMPUTER		AI AGENT
CPU Powerful but inert — no memory, no storage, no I/O	=	**LLM** One-shot text generator — answers then stops
RAM Fast but limited capacity	=	**Context Window** Fast but limited — requires compaction and management
Disk Storage Persistent, large, slower access	=	**Memory & Databases** Session persistence, RAG, workspace files
Device Drivers Hardware abstraction layer	=	**Tools & MCP** File I/O, bash, search, APIs, skills, integrations
Operating System Decides what the CPU sees and when	=	**Agent Harness** Decides what the LLM sees and when

DJIA Futures 51667.00 0.55% ↑ S&P 500 Futures 7421.25 0.26% ↑ Nasdaq F

THE WALL STREET JOURNAL.

SUBSCRIBE SIGN IN

NEWSLETTERS • CIO JOURNAL

Cisco Gave All 90,000 Employees Their Own AI Agent

The tech company granted access to a personal agent for all global employees this week

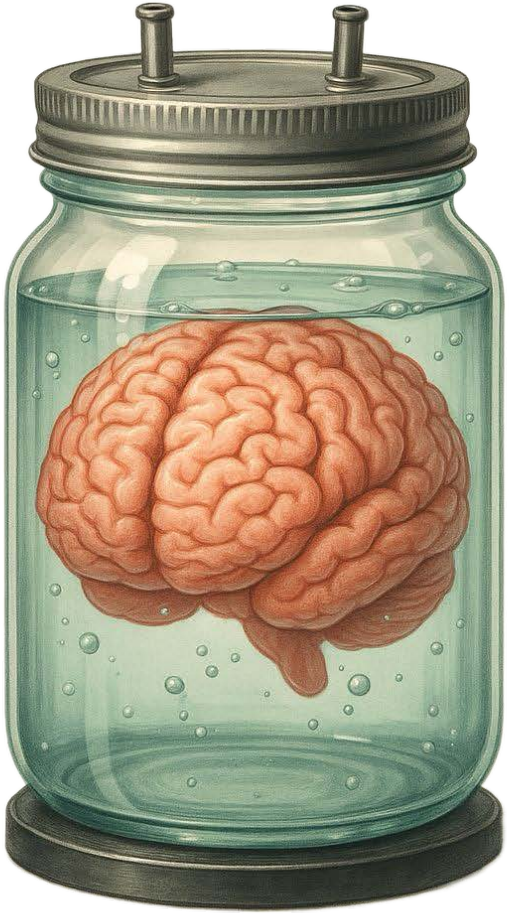


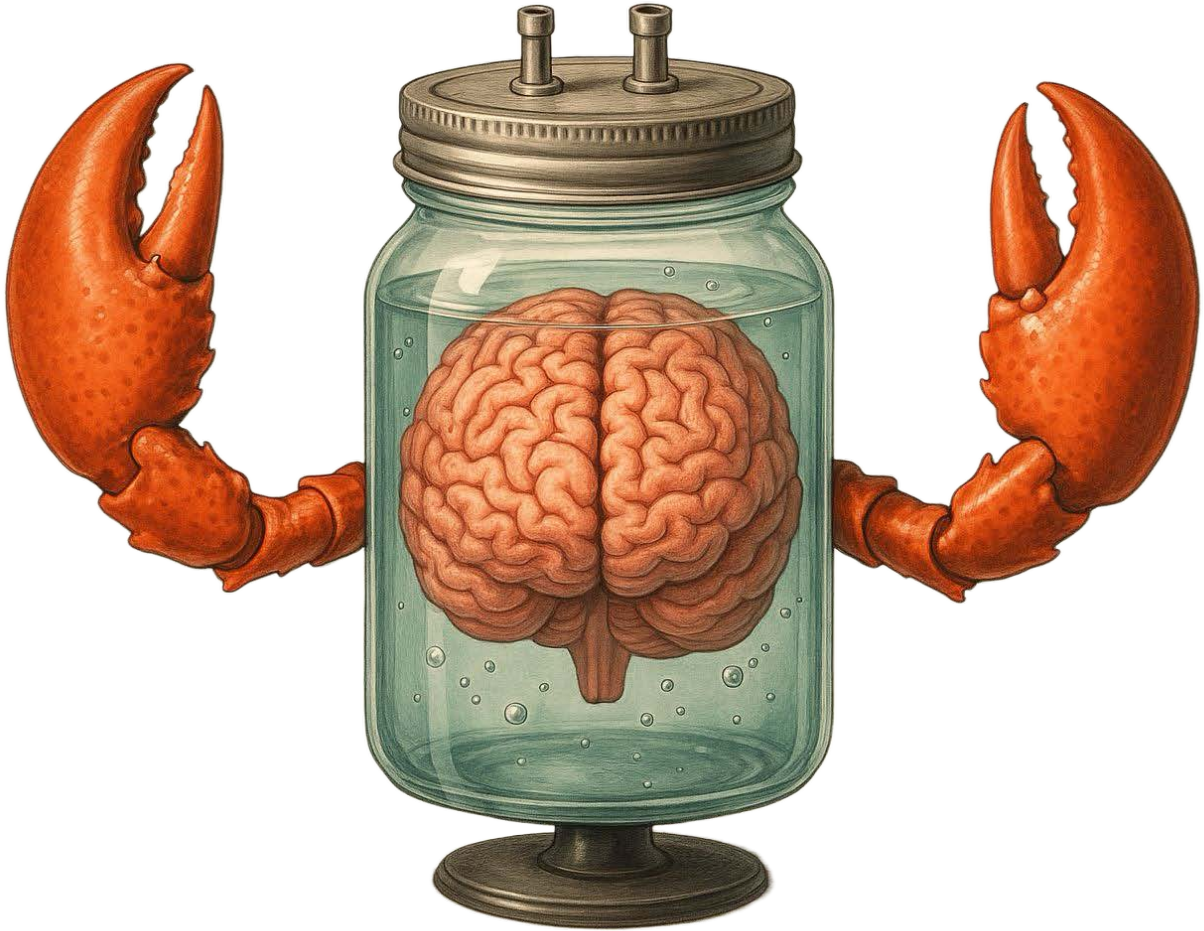
By Belle Lin
Aug. 27, 2026 9:09 am ET

Aa | Listen (1 min) ⋮









"Every company in the world needs an OpenClaw strategy. This is as big as HTML, as big as Linux."

"OpenClaw has open-sourced the operating system of agent computers. It's no different than how Windows made it possible for us to create personal computers. Now, OpenClaw has made it possible for us to create personal agents."

Jensen Huang on OpenClaw (GTC 2026, March 17, 2026)





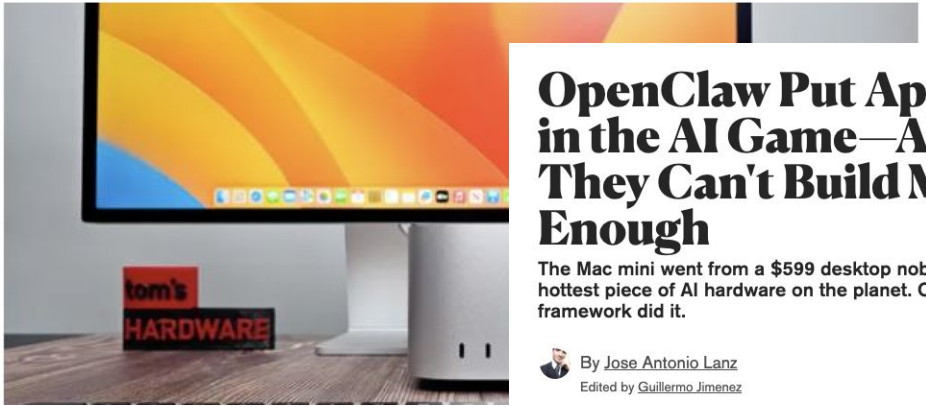
Fastest Growing OSS project in Github history

Project	Stars	Time to Reach
OpenClaw	369K	5 months (Nov 2025 → May 2026)
React	243K	~10 years
Linux	218K	~14 years
Kubernetes	100K	~3 years

OpenClaw-fueled ordering frenzy creates Apple Mac shortage — delivery for high Unified Memory units now ranges from 6 days to 6 weeks

News By Jowi Morales published February 15, 2026

AI is coming for high-end Mac Studios and Mac minis, too.



OpenClaw Put Apple Back in the AI Game—And Now They Can't Build Macs Fast Enough

The Mac mini went from a $599 desktop nobody cared about to the hottest piece of AI hardware on the planet. One open-source agent framework did it.

By Jose Antonio Lanz
Edited by Guillermo Jimenez

May 2, 2026
5 min read



Lobster buffet: China's tech firms feast on OpenClaw as companies race to deploy AI agents

PUBLISHED THU, MAR 12 2026•5:11 AM EDT | UPDATED THU, MAR 12 2026•7:31 PM EDT



China catches 'lobster fever' as cloud companies embrace OpenClaw

Ordinary Chinese face the harsh reality about the cost of "raising a lobster"

Written by
Oiwan Lam

Read this post in Español, عربي, русский

Posted 6 April 2026

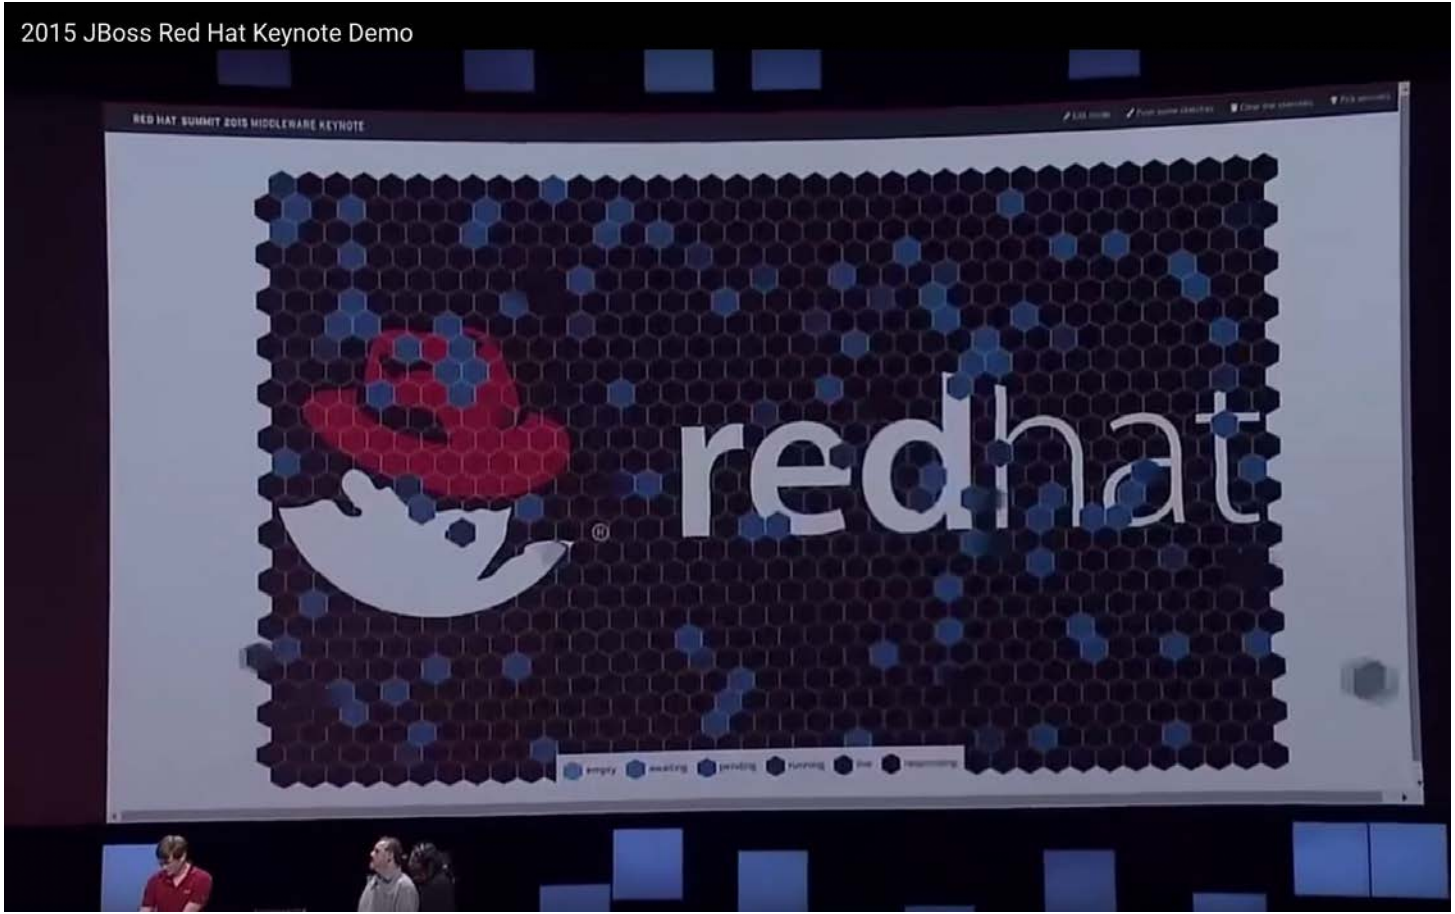




Bringing people closer to AI.

AI is changing the world faster than any technology in history. The closer you are to it, the less scary it feels. The OpenClaw Foundation exists to close that distance and make AI personal, fun, and empowering for everyone.

Get started →



https://www.youtube.com/watch?v=GCtpncA0Ea0

Red Hat Summit

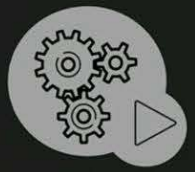
2015

Launch 1000+ Containers

Audience Claims a Container

Kubernetes Terms

Pod



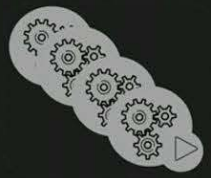
- ✓ 1+ containers
- ✓ Shared IP
- ✓ Shared storage volume
- ✓ Shared resources
- ✓ Shared lifecycle

Replication Controller/ Deployment



- ✓ The Desired State - replicas, pod template: health checks, resources, image

Service



- ✓ Grouping of pods (acting as one) has stable virtual IP and DNS name

Label



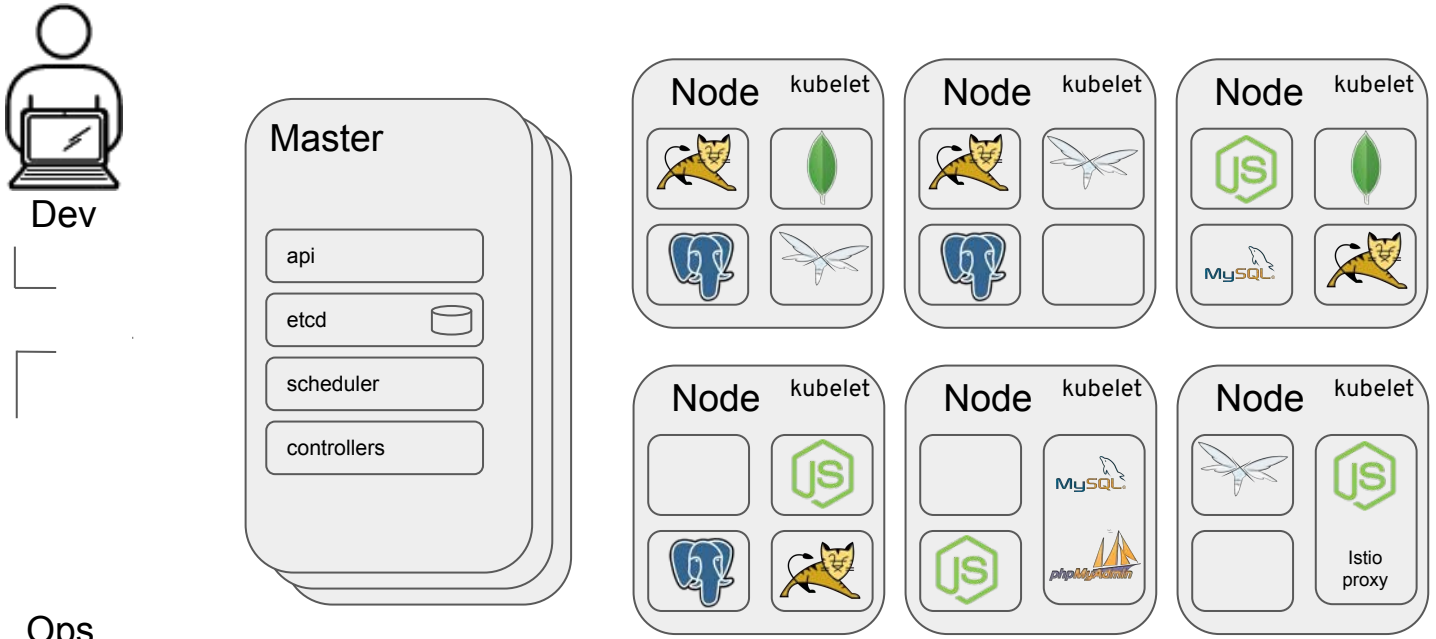
- ✓ Key/Value pairs associated with Kubernetes objects (env=production)

@burrsutter - bit.ly/9stepsawesome

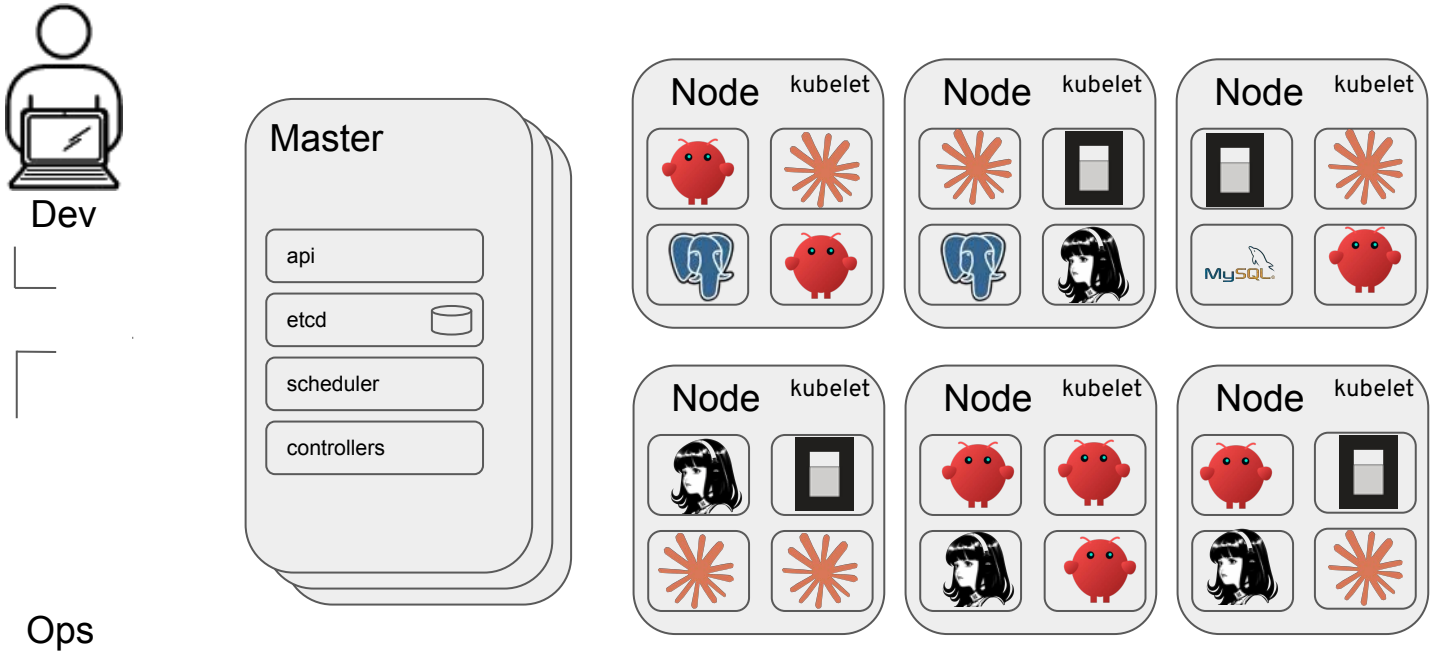
RED HAT DEVELOPER



Kubernetes Cluster



Kubernetes Cluster




Red Hat
OpenShift
OpenClaw pod-01
OpenClaw pod-02
OpenClaw pod-03
OpenClaw pod-04
OpenClaw pod-05
OpenClaw pod-06
OpenClaw pod-07
OpenClaw pod-08
OpenClaw pod-09
OpenClaw pod-10
OpenClaw pod-11
OpenClaw pod-12
OpenClaw pod-13
OpenClaw pod-14
OpenClaw pod-15
OpenClaw pod-16
OpenClaw pod-17
OpenClaw pod-18
OpenClaw pod-19
OpenClaw pod-20
OpenClaw pod-21
OpenClaw pod-22
OpenClaw pod-23
OpenClaw pod-24
OpenClaw pod-25
OpenClaw pod-26
OpenClaw pod-27
OpenClaw pod-28
OpenClaw pod-29
OpenClaw pod-30
OpenClaw pod-31
OpenClaw pod-32
OpenClaw pod-33
OpenClaw pod-34
OpenClaw pod-35
OpenClaw pod-36
OpenClaw pod-37
OpenClaw pod-38
OpenClaw pod-39
OpenClaw pod-40
OpenClaw pod-41
OpenClaw pod-42
OpenClaw pod-43
OpenClaw pod-44
OpenClaw pod-45
OpenClaw pod-46
OpenClaw pod-47
OpenClaw pod-48

Isolate the Agent

What you can do today (with no new code required)

Red Hat Enterprise Linux

RHEL

OS & Container Isolation

- SELinux
- Seccomp
- Rootless Podman

Red Hat OpenShift

OpenShift

Platform & Network Isolation

- SCC (restricted-v2)
- Network Policy
- RBAC / Admission Policy

Red Hat Ansible Automation Platform

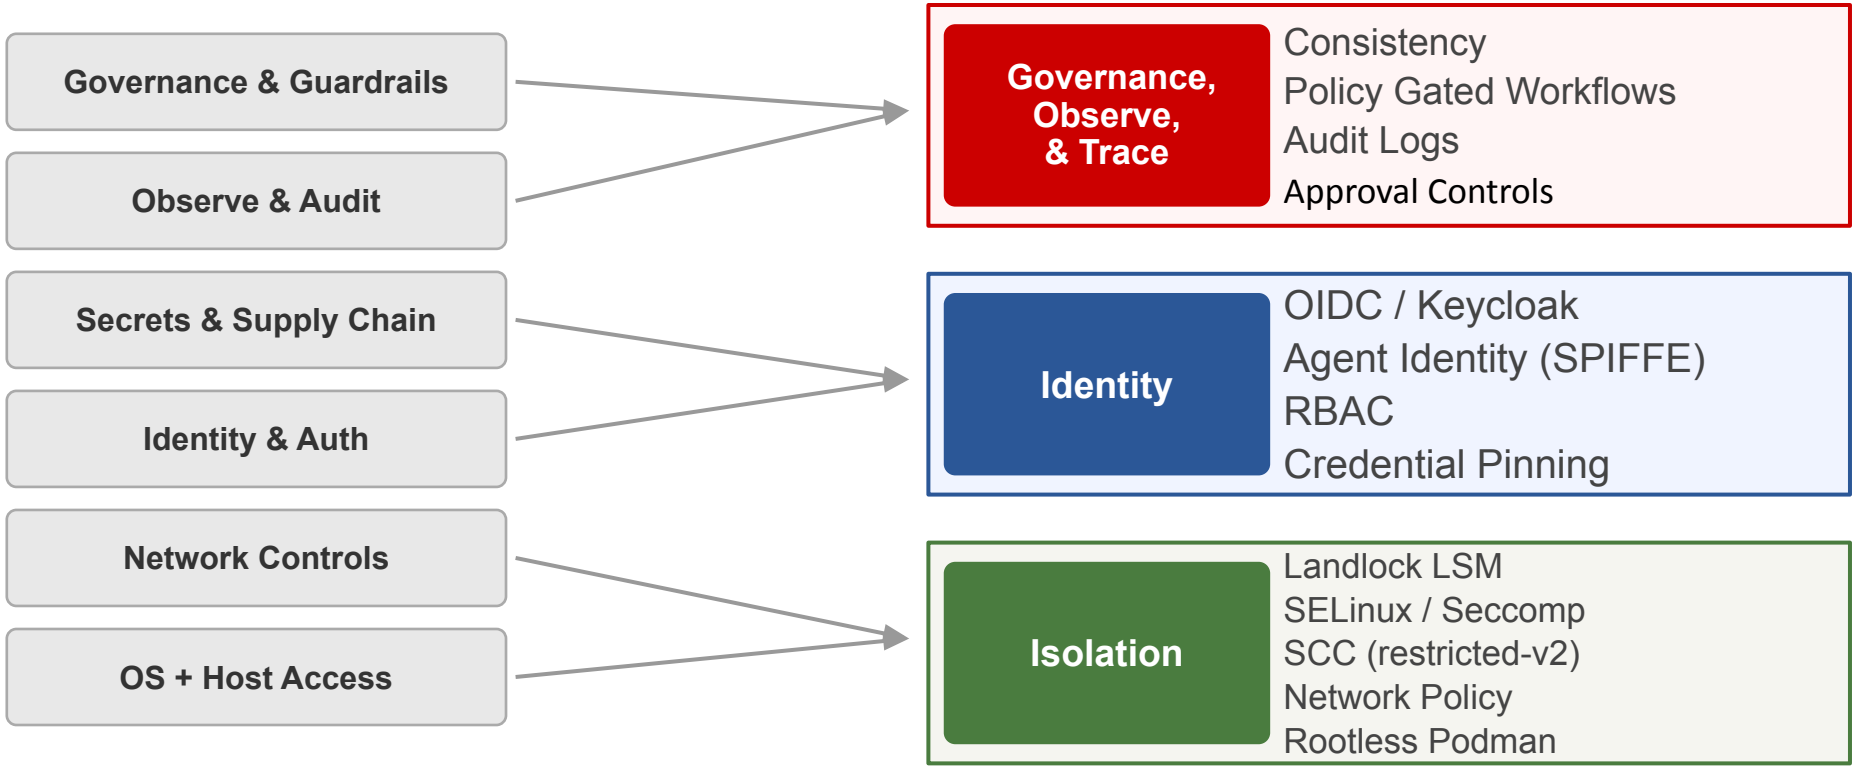
Ansible

Automation & Governance

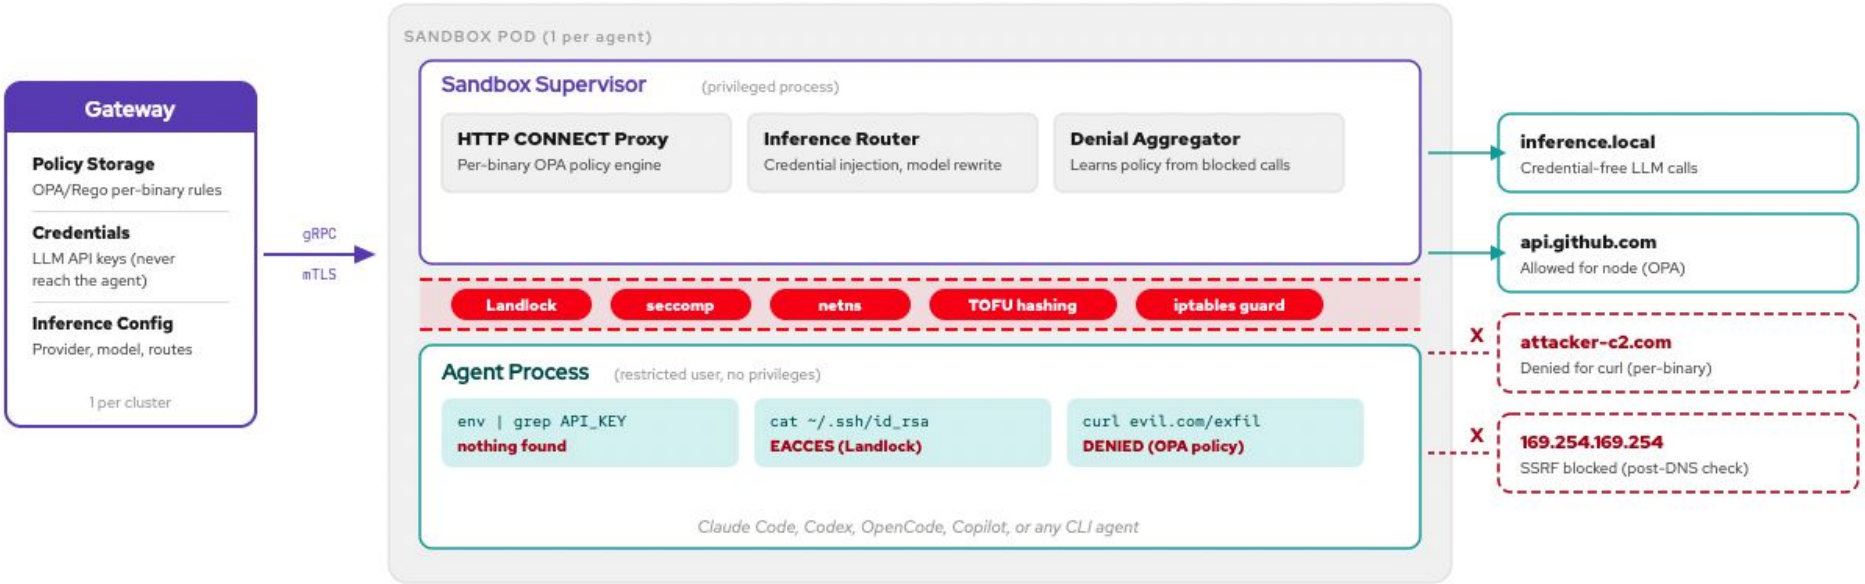
- Policy-Gated Workflows
- Approval Controls
- Execution Audit Trail

AgentOps Security for the Future

Defense-in-Depth for AI Workloads



OpenShell (Software Sandboxing)





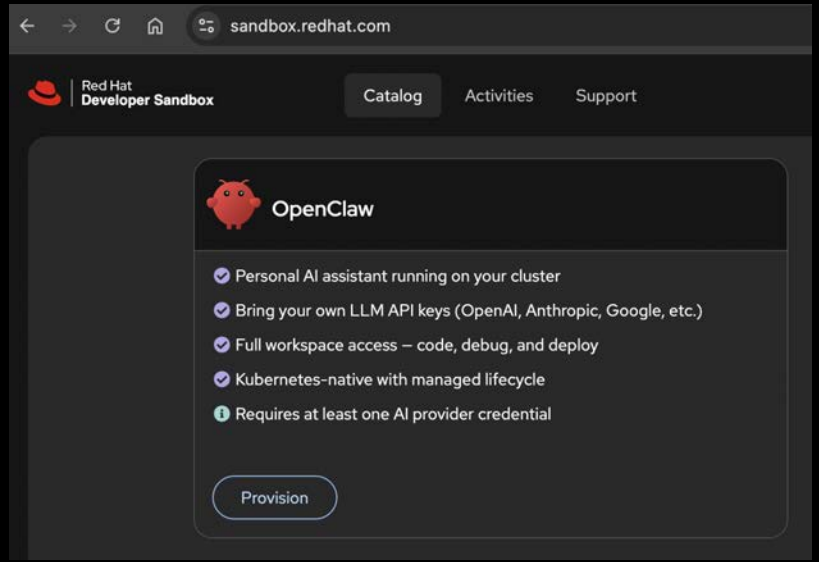
It’s here

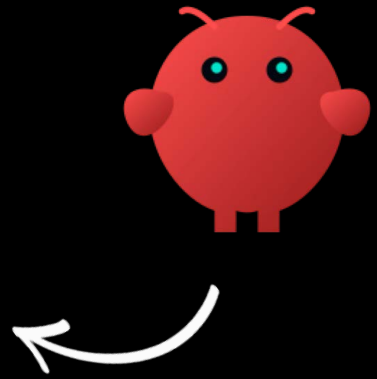
OpenShift 5 is GA!

red.ht/yougetaclaw


From the deepest terror comes the deadliest catch.
CLAWS
STARRING
PAT O'REILLY · JOHN HAYES · MARINA DONOVAN
DIRECTED BY
BURR SUTTER
WRITTEN BY
SALLY O'MALLEY
PRODUCED BY
IT PRODUCTIONS
R
RESTRICTED
UNDER 17 REQUIRES ACCOMPANYING PARENT OR ADULT GUARDIAN
IT PRODUCTIONS
A BOSTON, MA PICTURE
COLOR BY OCEANVISION

Launch an OpenClaw on OpenShift today!





developers.redhat.com/developer-sandbox