

citrix®

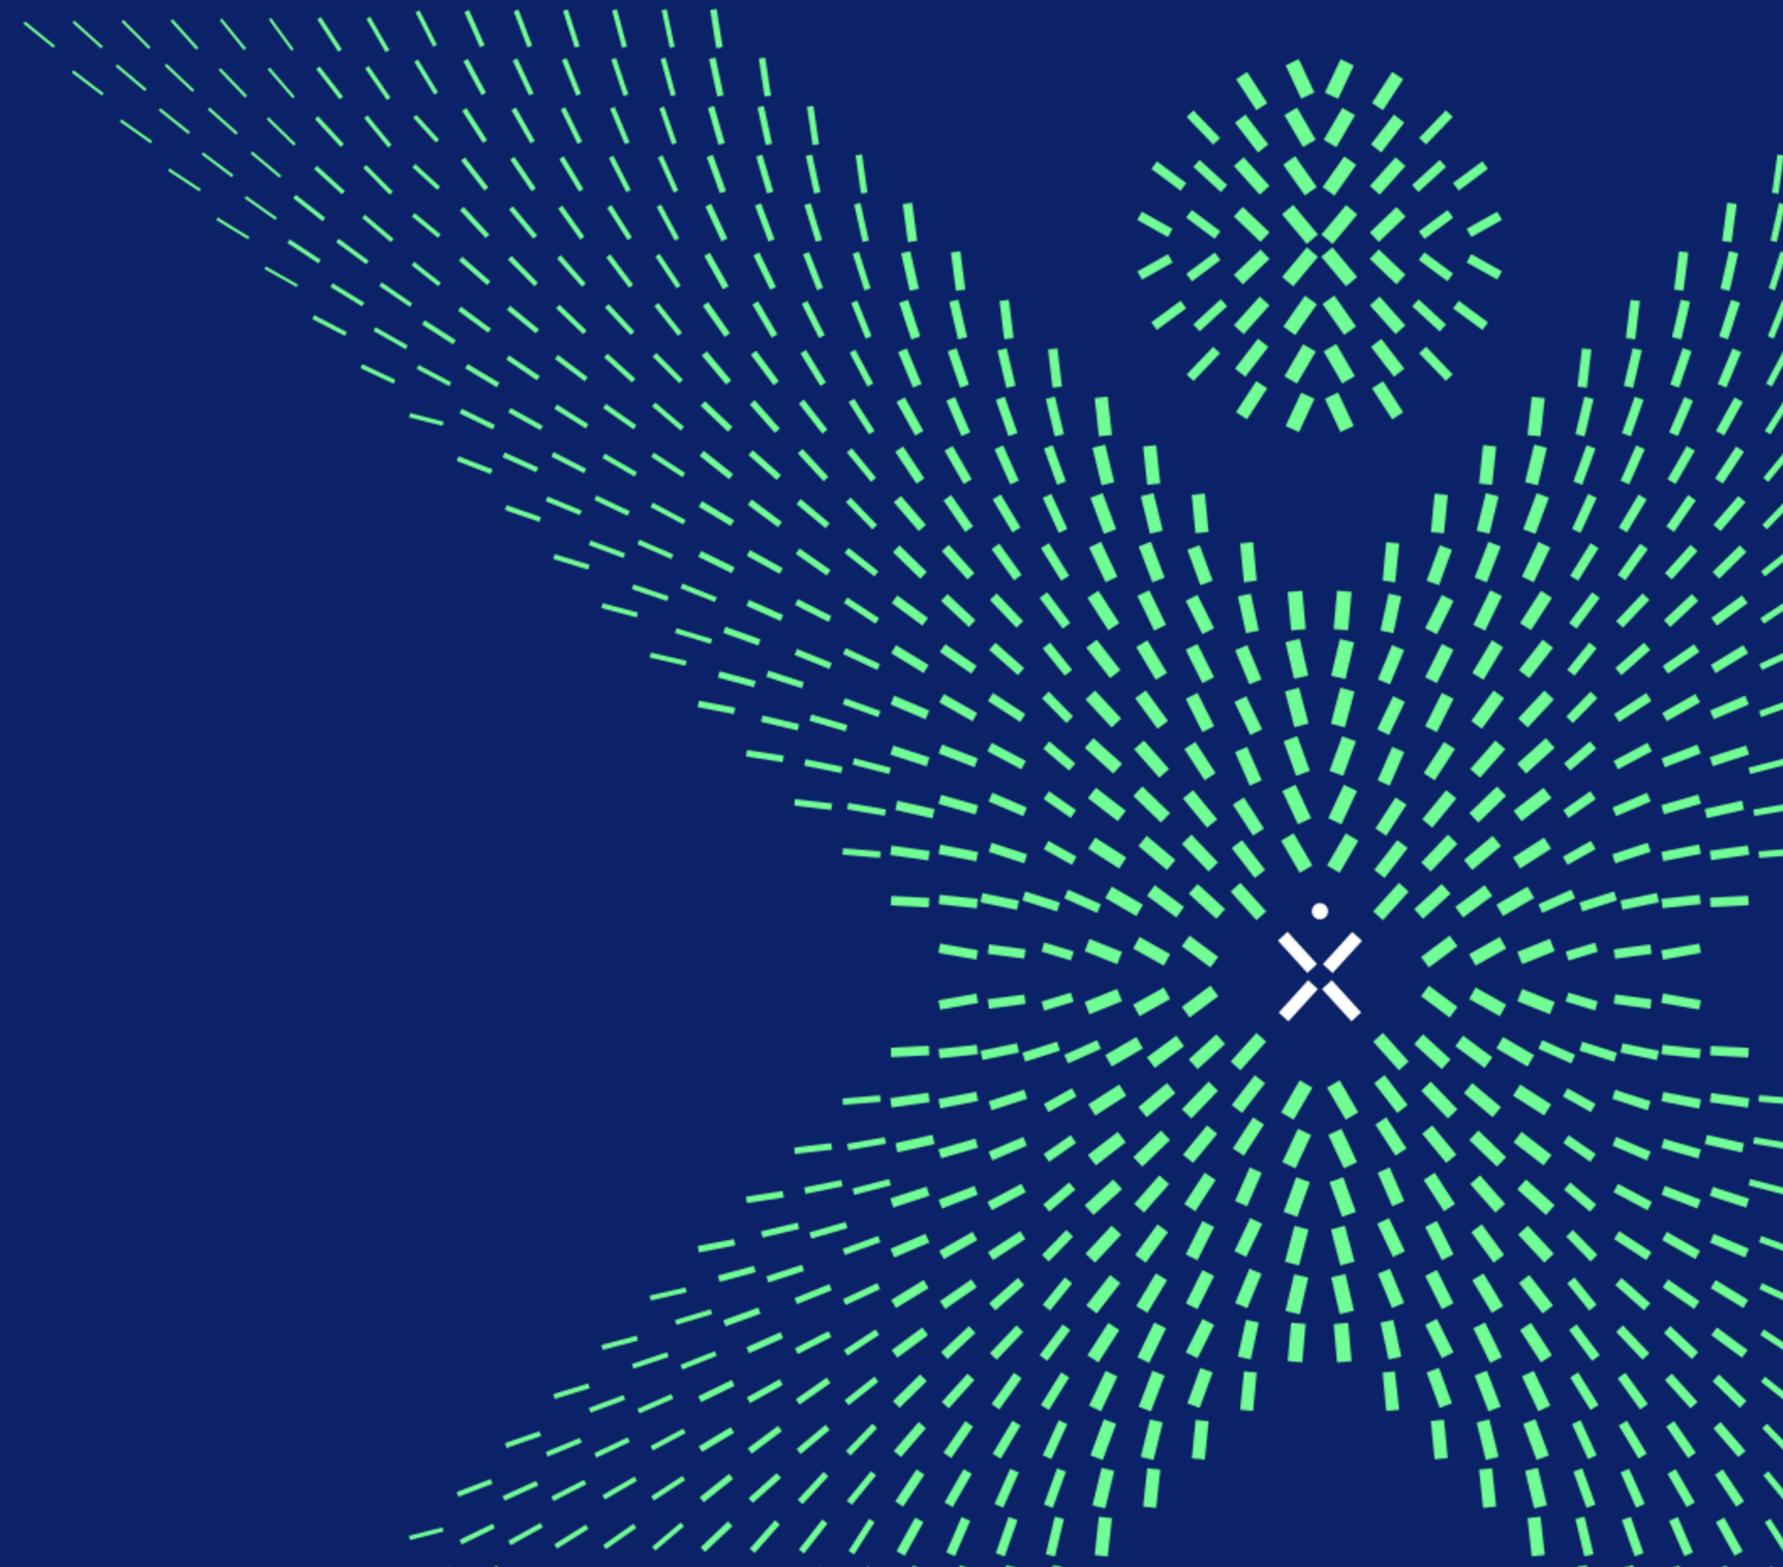


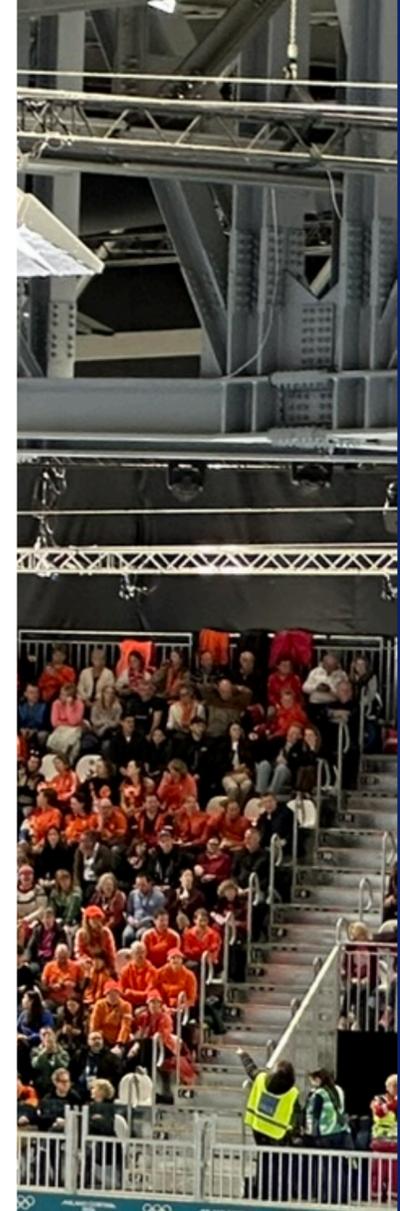
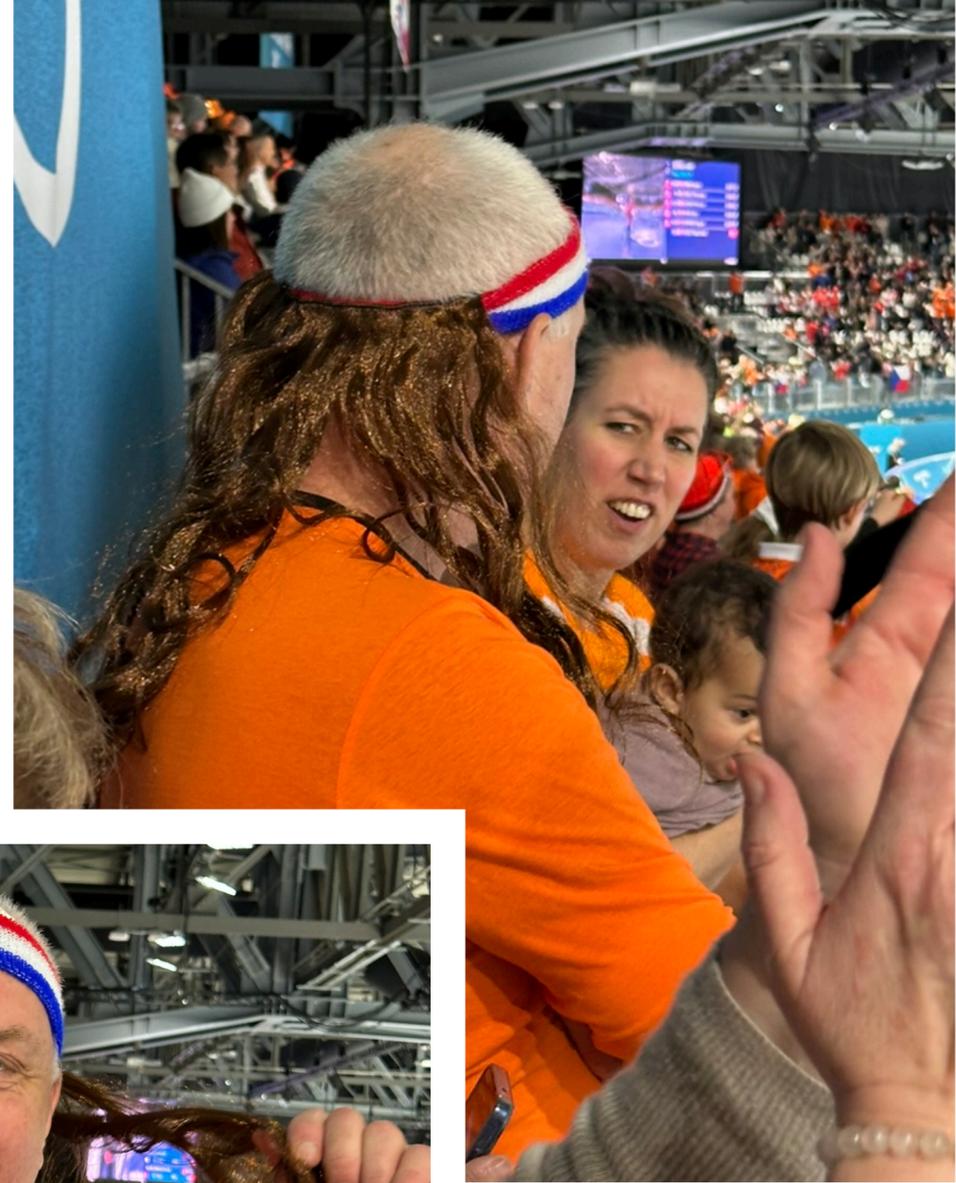
DeCUG #28

The New Cognitive Stack

Brian Madden
Futurist, Citrix

18 March 2026





“Nearly eight in ten companies report using gen AI – yet just as many report no significant bottom-line impact.”

Companies Are Pouring Billions Into A.I. It Has Yet to Pay Off.

Corporate spending on artificial intelligence is surging as executives bank on major efficiency gains. So far, they report little effect to the bottom line.

Listen to this article · 8:36 min [Learn more](#) Share full article



By Steve Lohr
Steve Lohr has written about technology and its effect on business and the economy more than 20 years.

Aug. 13, 2025
[Leer en español](#)

Nearly four decades ago, when the personal computer boom was in full swing, a phenomenon known as the “productivity paradox” emerged.

It was a reference to how, despite companies’ huge investments in new technology, there was scant evidence of a corresponding gain in work efficiency.

Today, the same paradox is appearing, but with generative artificial intelligence. According to [recent research from McKinsey & Company](#), nearly eight in 10 companies have reported using generative A.I., but as many have reported “no significant bottom-line impact.”

A.I. [technology has been racing ahead](#) with [chatbots like ChatGPT](#), fueled by a high-stakes arms race among tech giants and superrich start-ups and prompting an expectation that everything from back-office accounting to

NEWSLETTERS · CFO DAILY

MIT report: 95% of generative AI pilots at companies are failing

BY SHERYL ESTRADA
SENIOR WRITER AND AUTHOR OF CFO DAILY

August 18, 2025 at 6:54 AM EDT

Seizing the



Most knowledge work is invisible

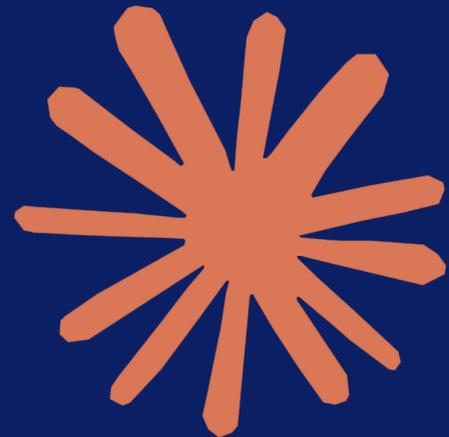
emails
documents
meeting transcripts
chats



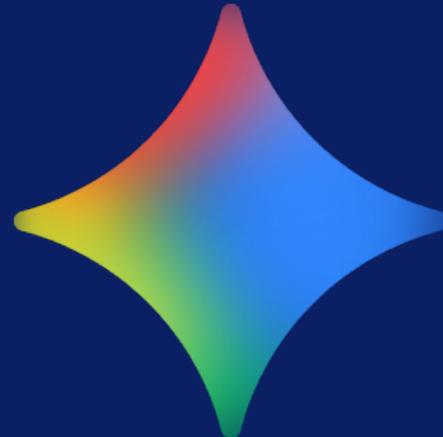
thinking
reasoning
staring out the window
looking at birds

Something happened in November 2025

(4 months ago)



Opus 4.5



Gemini 3



GPT 5.2

Something happened in November 2025

“AI helps me code”



Opus 4.5

Boris Cherny

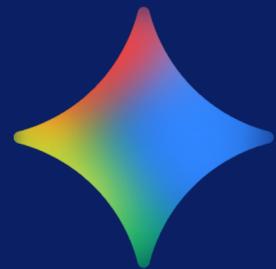
“100% of my contributions were written by Claude Code”

Brian Madden

“HOLY
F★★★★★
S★★★★!!”

Andrej Karpathy

“There has been a phase change”



Gemini 3

Paul Ford

“I spent an entire session of therapy talking about it”

Simon Willison

“It’s all changed since November”



GPT 5.2

OCT 2025

NOV 2025

DEC 2025





What will knowledge work be in 18 months? Look at what AI is doing to coding right now.



Brian Madden



There's a lot of buzz about **something big happening** in software engineering thanks to the latest batch of AI models. However, most knowledge workers think this is just a "coding thing" which doesn't apply to them. They're wrong.

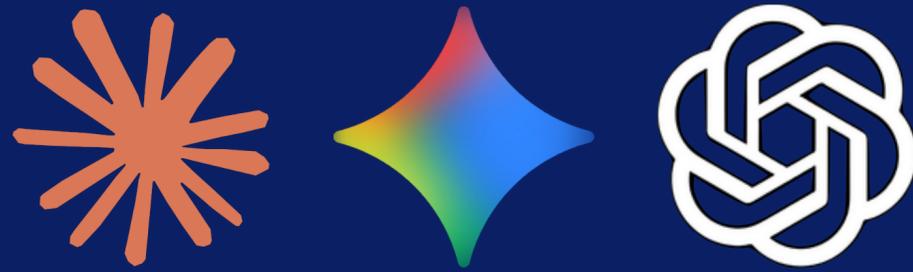
Nate B. Jones

Writing on how software engineers need to change

knowledge workers

Most ~~engineers~~ have spent years developing their intuitions around **writing work docs** ~~implementation~~ and those are now not super useful.

The new skill is describing what you want precisely enough that AI can **goals & objectives** build it, and then writing ~~tests~~ that capture what you actually need, and **ideas & docs** reviewing AI-generated ~~code~~ for subtle conceptual errors rather than **formatting** simple ~~syntax~~ mistakes.



November 2025

Minutes of focus → Hours of sustained work

Loses track after several tasks → Maintains its own todo lists

Hilarious spirals on errors → Diagnoses and recovers

Tool use mostly works → Dependable agentic work

“Technically” works → Seriously works

OPINION
GUEST ESSAY

The A.I. Disruption We've Been Waiting for Has Arrived



THE 2028 GLOBAL INTELLIGENCE CRISIS

A Thought Exercise in Financial History, from the Future

CITRINI AND ALAP SHAH
FEB 22, 2026

Something Big Is Happening

6 k 41 k 118 k 86 M

Think back to February 2020.

If you were paying close attention, you might have noticed a few people

talking about a v
close attention. T
school, you were
trips. If someone
have thought the
the internet. The
world changed. Y
rearranged itself

Anthropic's 'Claude Cowork' Release Triggers \$285 Billion 'SaaSocalypse': A Brutal Wake-Up Call for Legacy Tech and Finance

By: [MarketMinute](#)
February 06, 2026 at 11:41 AM EST

TECHNOLOGY

AMERICA ISN'T READY FOR WHAT AI WILL DO TO JOBS

Does anyone have a plan for what happens next?

By Josh Tyrangiel

Illustrations by Stephan Dybus

LIFESTYLE • CAREERS

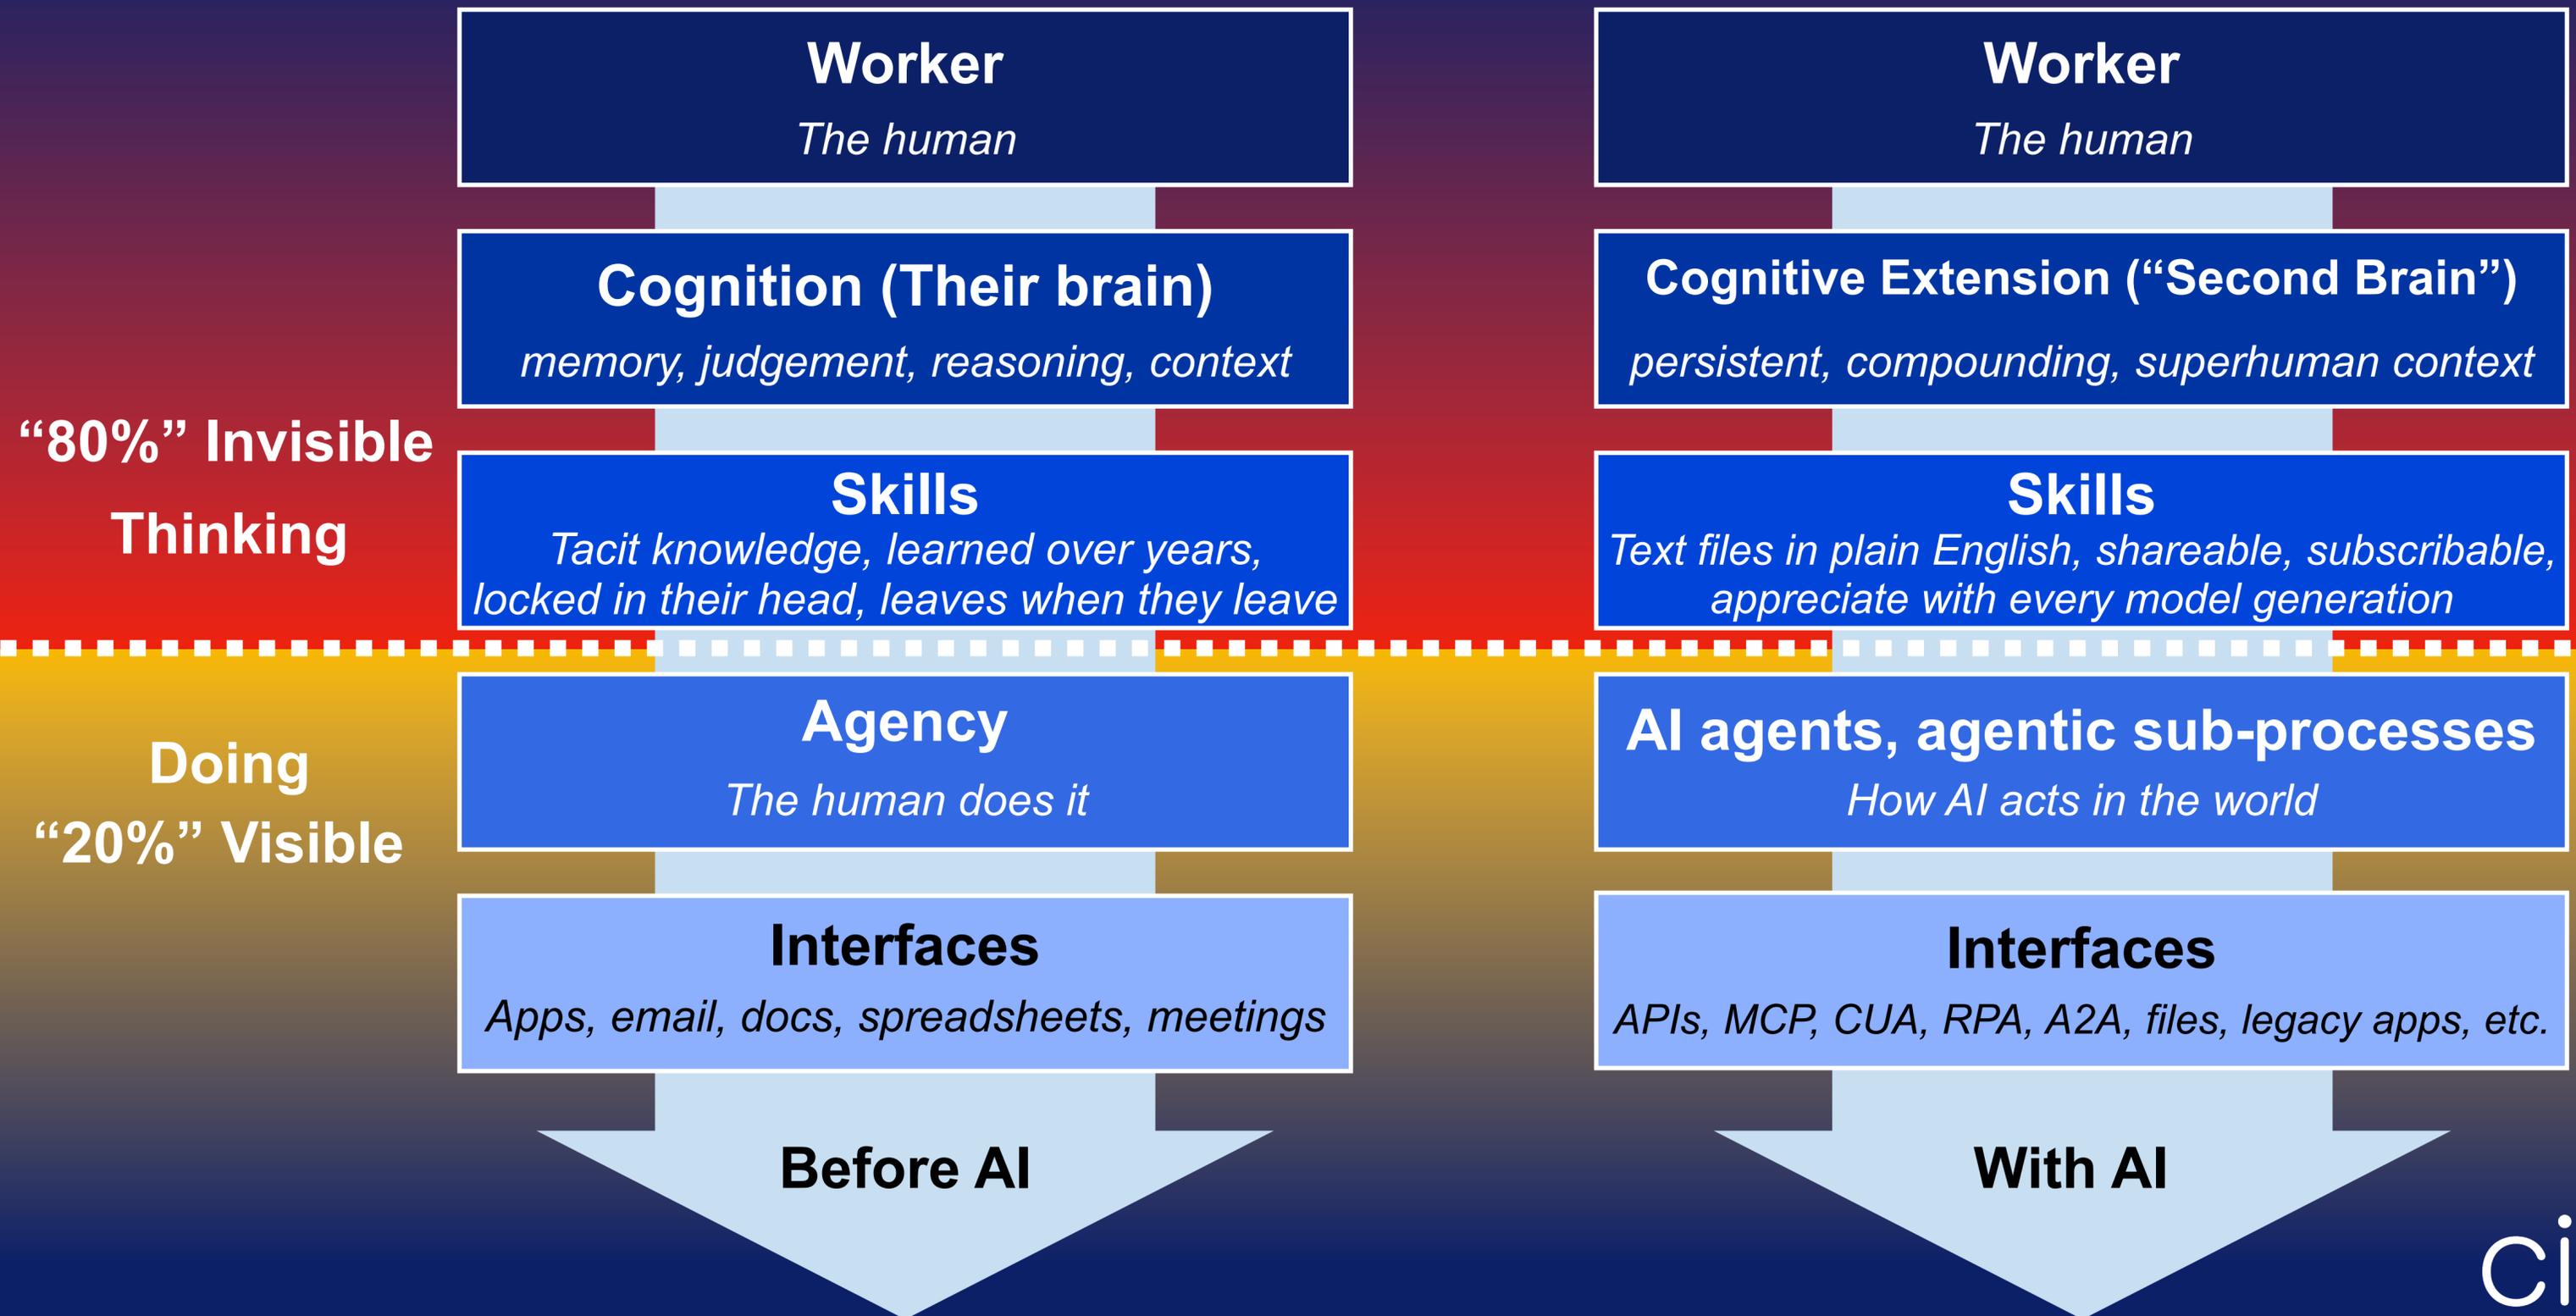
Workers Are Afraid AI Will Take Their Jobs. They're Missing the Bigger Danger.

It isn't whether artificial intelligence is going to replace them. It's who will control the knowledge that companies capture from their employees.

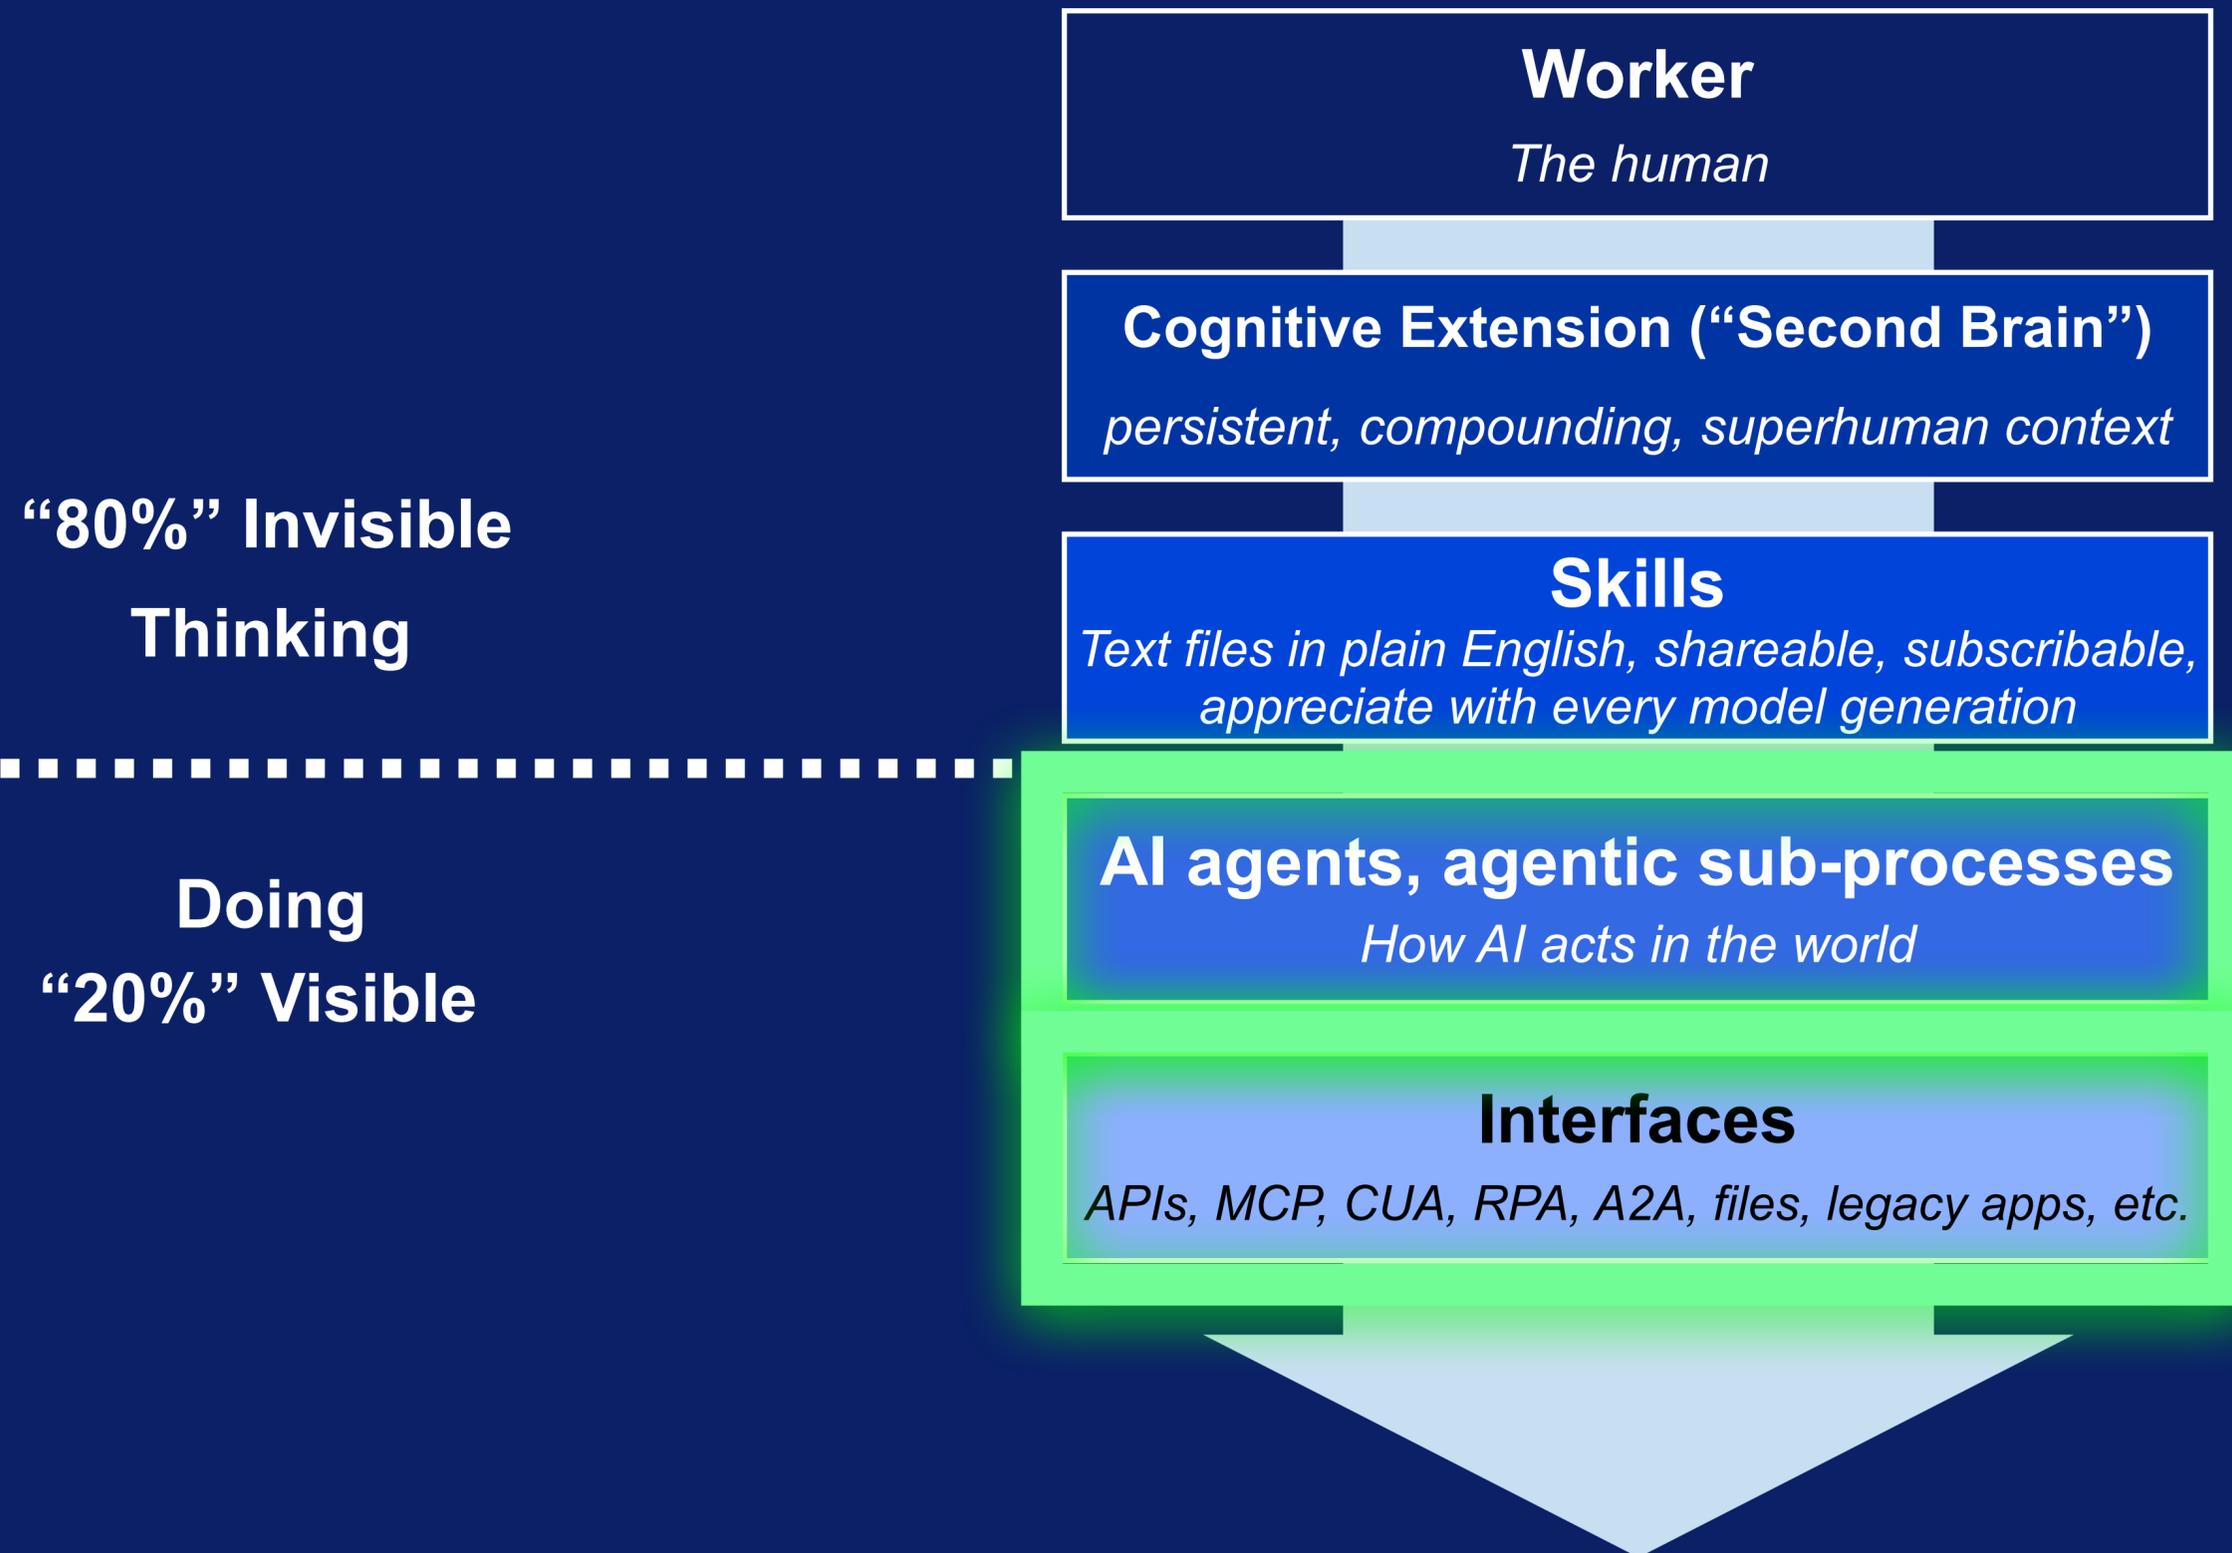
By [Matthew Call](#)
Feb. 15, 2026 at 12:00 pm ET



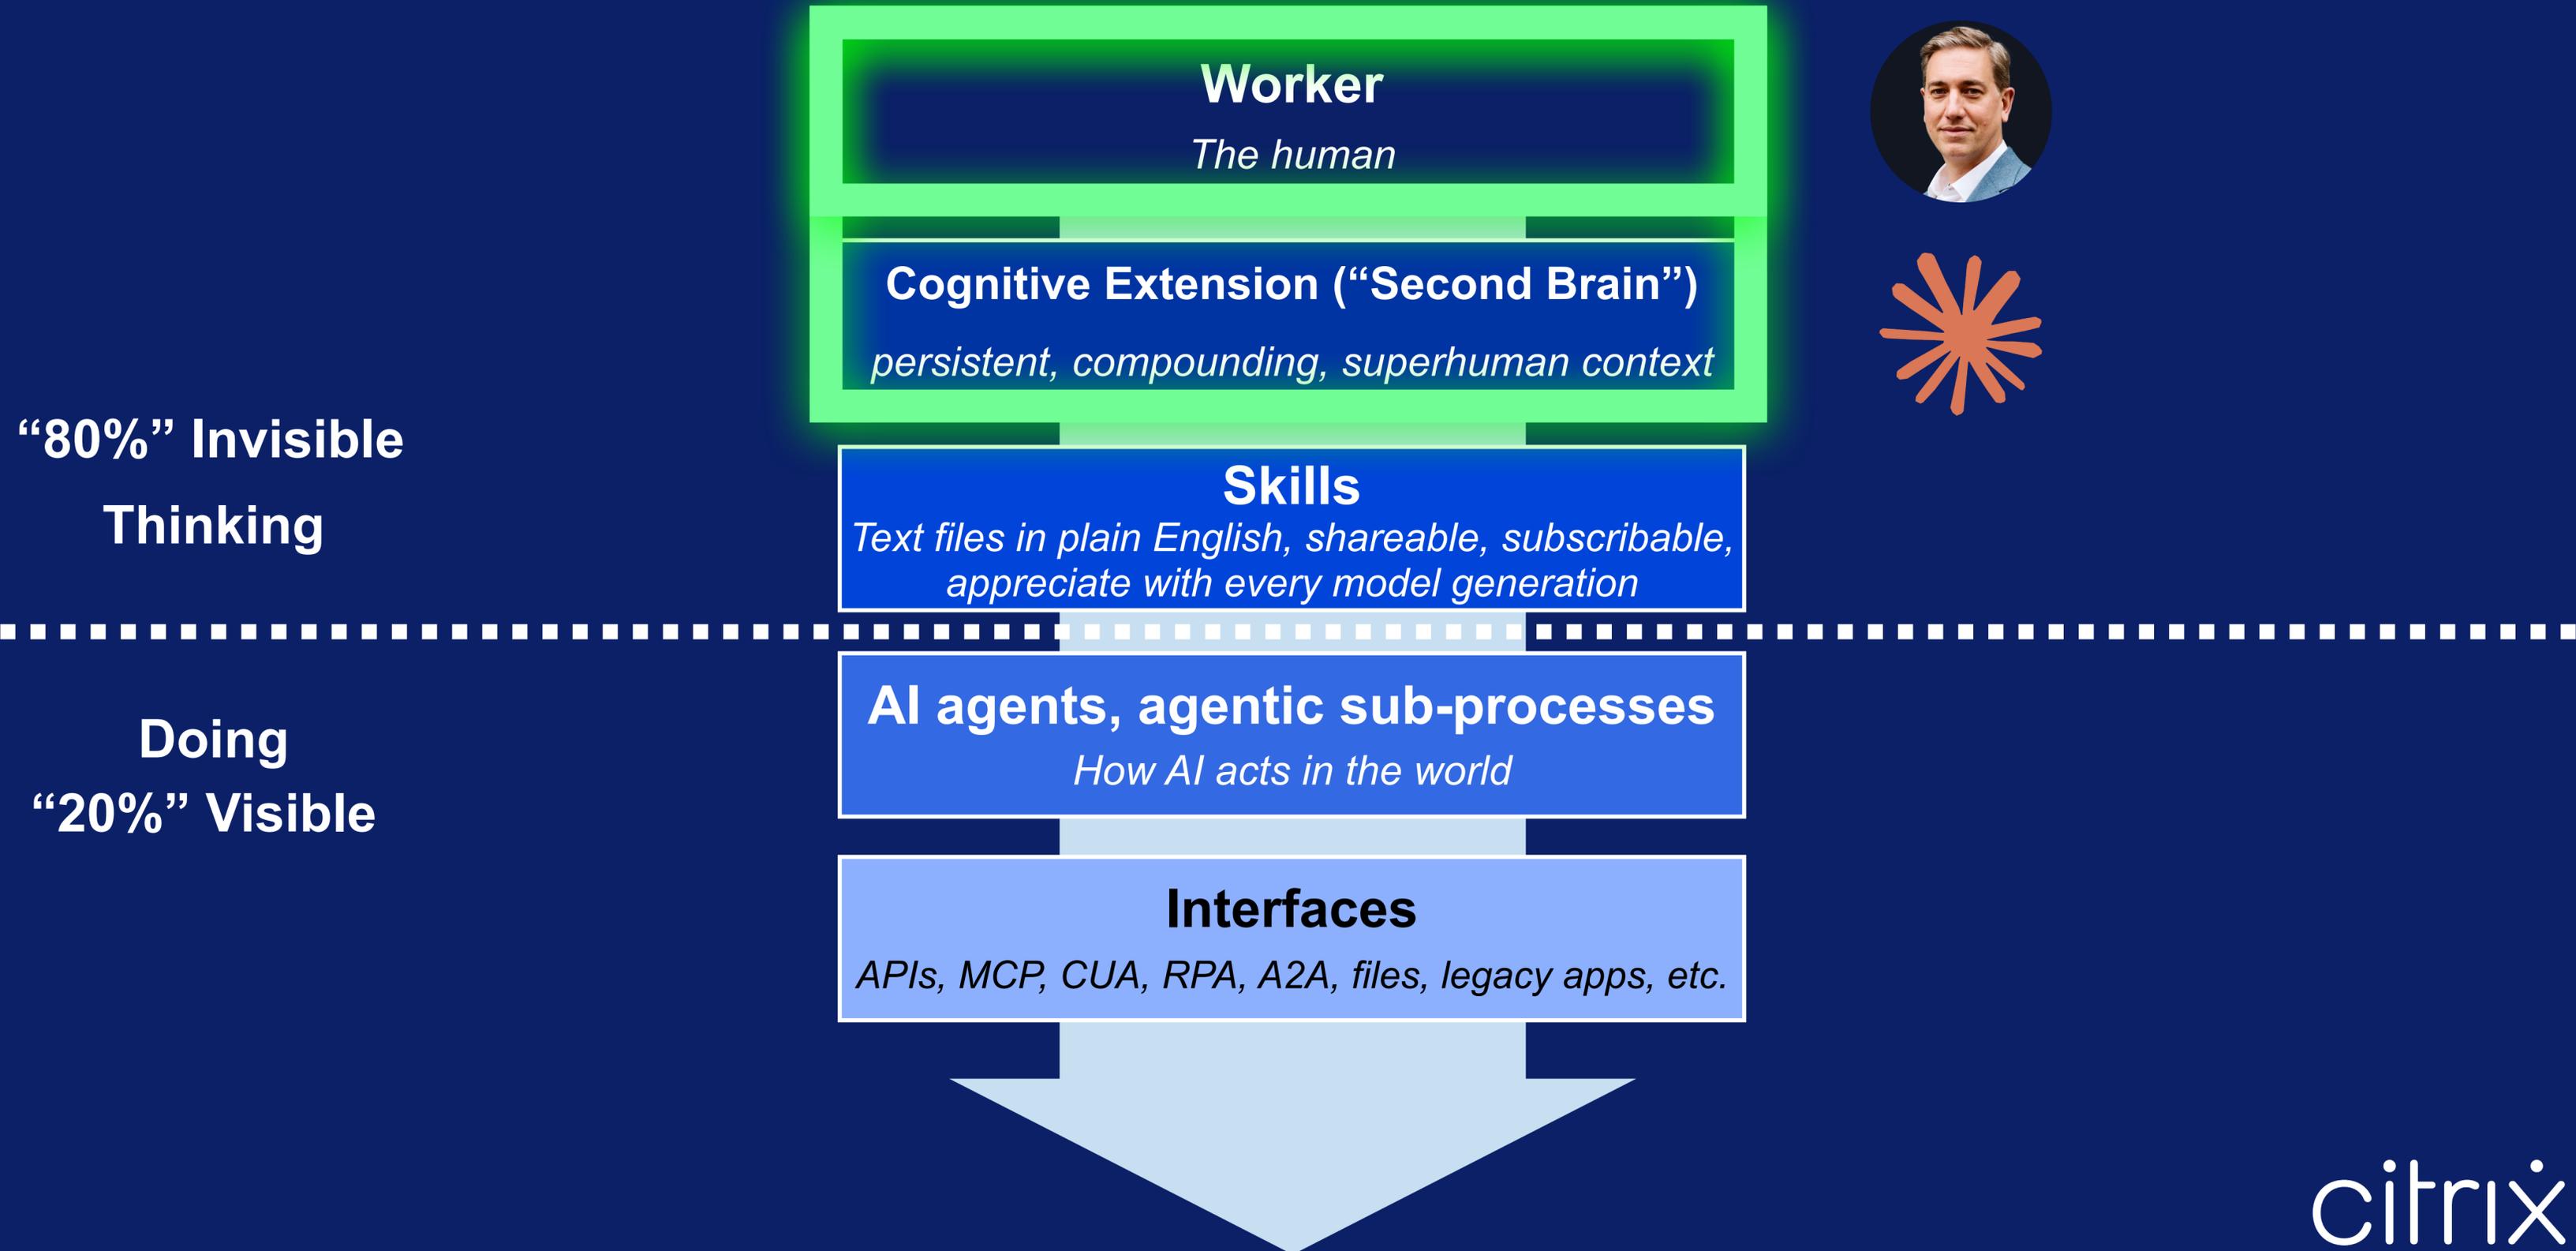
The cognitive stack of knowledge work

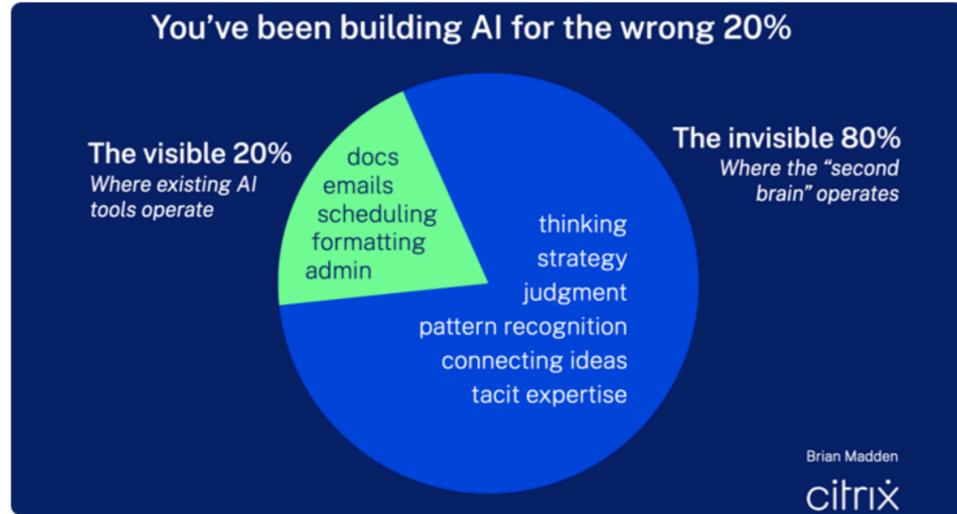


The cognitive stack of AI-enhanced knowledge work



The cognitive stack of AI-enhanced knowledge work





I built a second brain using AI, and it's changed the way I work. THIS is the future of knowledge work.

Brian Madden
Publié le 10 févr. 2026

+ Suivre

I've spent over 30 years in enterprise IT. I've led analyst teams, done corporate strategy, built & sold companies, and now I'm VP & Futurist at Citrix. All of that knowledge—every framework, every hard-won insight, every pattern I've learned to recognize—lived in one place: my head. If you wanted access to it, you had to talk to me. And even then, I could only leverage what I happened to remember in the moment.

That all changed a few weeks ago. Now my knowledge lives in a folder of plain text files on my laptop. An AI reads them, maintains them, and updates them based on everything I do, continuously, every single day.

That simple concept is a "second brain", and has so fundamentally changed the way I work that I now understand how this concept, when used by all workers, will profoundly change the way knowledge work happens. In fact all my thinking about the very near future of work has just blown up, which is what I'm going to walk through in this article.



Image by Gemini

Hey creators, stop publishing content. Start publishing your second brain.

Brian Madden
Publié le 17 févr. 2026

+ Suivre

Hey creators and experts: I will pay to subscribe to your brain.

UPDATE Feb 26: I just did this! [Your AI can now connect into my second brain!](#)

I don't mean your substack, I mean I'll pay for a feed of full text transcripts of everything you publish, maybe along with your other thoughts and opinions, which I'll merge into my own second brain so I can more easily incorporate your perspectives & thinking into my personal AI knowledge system. Since you're already putting the content out there in public, it's already possible for anyone to incorporate. I just want to make it easier, and pay you for it.

Quick catch-up: I recently [built a personal AI knowledge system which I call my "second brain"](#) that's essentially a folder of text files on my laptop that an AI reads, maintains, and compounds daily. It includes all my frameworks, positions, and publicly published materials, as well as my accumulated thinking from 32 years in enterprise IT. It's all structured, queryable, and building on itself every day. (And it's all maintained by AI, which has recently gotten really good at this kind of thing.)



Announcing brianmadden.ai : I just published my brain for you to merge into your AI

Brian Madden
Publié le 26 févr. 2026

+ Suivre

A few weeks ago I wrote about how I used AI to [build a personal compounding knowledge system](#), (what some call a "second brain"), which fundamentally changed how I work. Then last week I suggested that [experts should stop publishing content and instead publish their knowledge](#) as modules anyone can plug into their own AI.

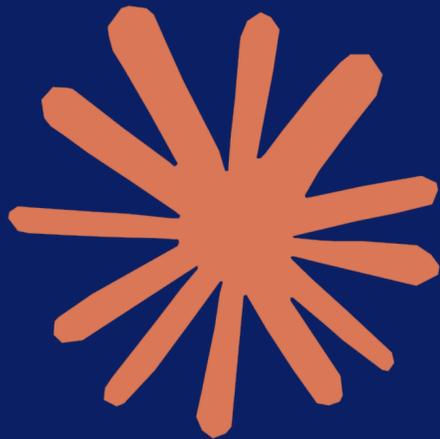
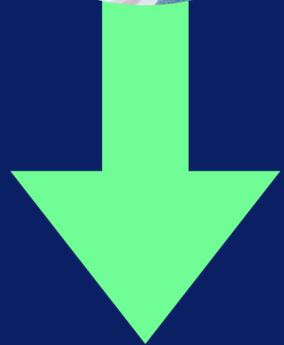
Today I'm happy to announce that I'm doing that, at [brianmadden.ai](#), which is open source and live today!

The brianmadden.ai repo contains all my published thinking on AI and the future of knowledge work based on my 32 years of expertise: all my Citrix blogs, LinkedIn posts, podcasts & speech transcripts, 9 standalone frameworks, and a 30,000-word synthesis that connects everything I've written. It's all structured so any AI assistant can load it and reason with my perspective.

So the "audience" for this is your AI system, not you. You can connect to it right now. It takes about 10 seconds.



AI as my cognitive extension



Chat Cwork Code

Write weekly blog post Preview

Let's work on this week's blog post.

Let me load up the relevant context.

Here's where things stand:

Last published: "Skills are all you need" (March 12)

Overdue in the editorial calendar:

1. "The bitter lesson, revisited" -- framework complete, target was week of March 3
3. "Enterprise invariants" -- idea complete, target was March 10-17 window

Fresh seeds from today: AI stack value layers, FDE companies as consulting firms

Which post do you want to work on? The bitter lesson and five levels are both framework-complete and just need drafting. Enterprise invariants is also ready. Or if something else is pulling at you right now.

main ← main +58 -1 Create PR

Reply...

+ </> Auto accept edits Opus 4.6

~/.../com-apple-CloudDocs/bmad

AI as my cognitive extension



READ ME FIRST

This is Brian Madden's personal AI operating system. Everything Claude needs to work with Brian is contained in this folder structure.

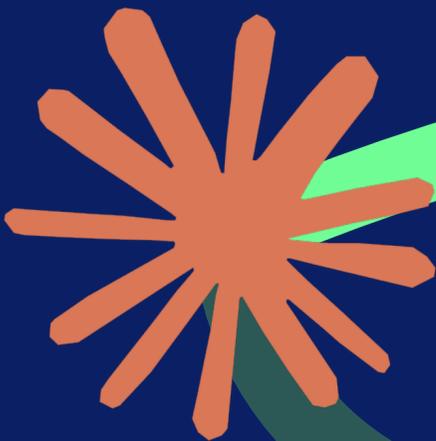
For Claude: How This Works

On every session start:

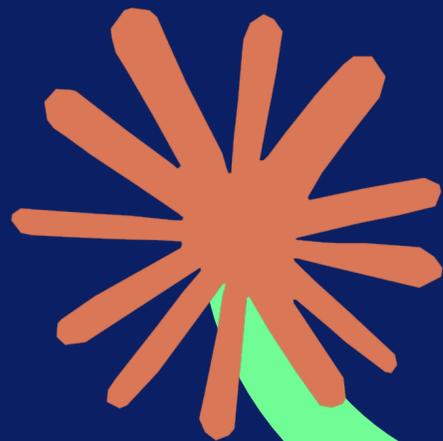
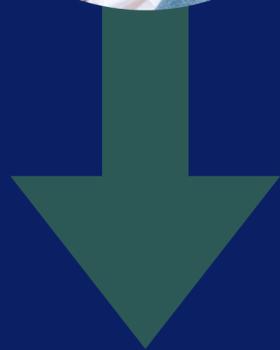
1. Read this file first
2. Read `[[index]]` to understand what's where
3. Read `[[me/thinking]]` to know where Brian's head is right now
4. Read `[[me/preferences]]` for behavior rules
5. Read `[[skills/index]]` to know available skills and triggers
6. Read `[[skills/lessons]]` for mistake patterns to avoid
7. Load other files as relevant to the current task

Core principle: This system is self-contained. Do not rely on external memory or past conversations. Everything you need is in these files.

File index:



AI as my cognitive extension



THINKING

Where my thinking is right now

This is the frontier map. [[me/synthesis]] captures what I've published. This captures where my head is **today**—the arguments forming, the connections emerging, the questions I'm chewing on. Updated frequently. Read this to know where I'm heading, not just where I've been.

****Last updated:**** March 17, 2026

The big arguments I'm developing

Second brains as infrastructure, not productivity hack
The published anchor article and Citrix blog established the concept. What's forming now is the **infrastructure** argument: second brains aren't a personal productivity trick, they're the individual-scale version of a new enterprise layer. Context graphs (Foundation Capital thesis) are the enterprise version. Subscribable brains are the distribution model. The 80/20 framework holds across all three levels—individual

AI as my cognitive extension



READ ME FIRST

This is Brian Madden's personal AI operating system. Everything Claude needs to work with Brian is contained in this folder structure.

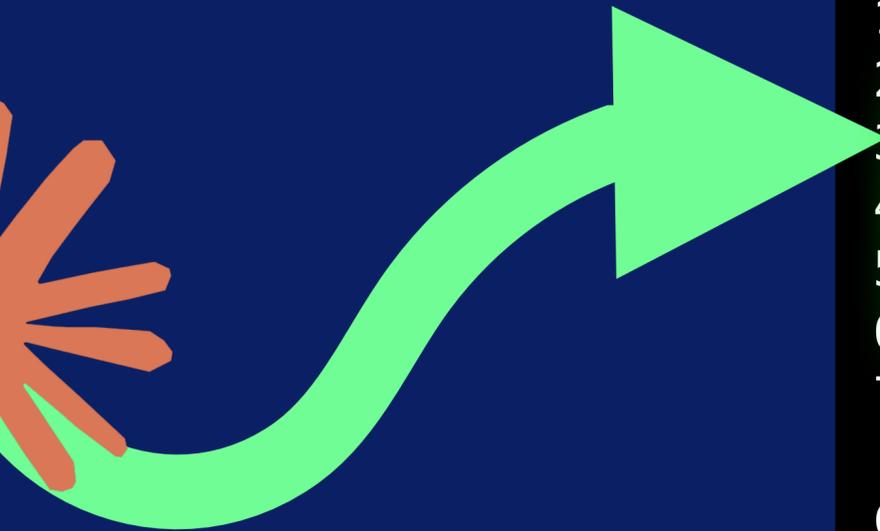
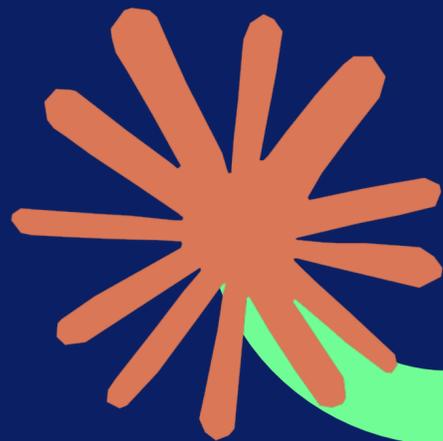
For Claude: How This Works

On every session start:

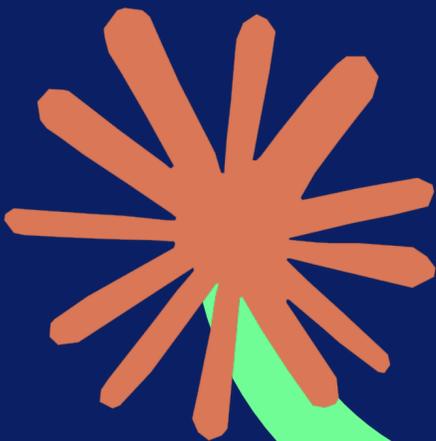
1. Read this file first
2. Read `[[index]]` to understand what's where
3. Read `[[me/thinking]]` to know where Brian's head is right now
4. Read `[[rules]]` for behavior rules
5. Read `[[skills/index]]` to know available skills and triggers
6. Read `[[skills/lessons]]` for mistake patterns to avoid
7. Load other files as relevant to the current task

Core principle: This system is self-contained. Do not rely on external memory or past conversations. Everything you need is in these files.

File index:



AI as my cognitive extension



SKILLS INDEX

```
# Skills Index
```

```
Skills are procedures Claude can follow. When Brian asks to do something
```

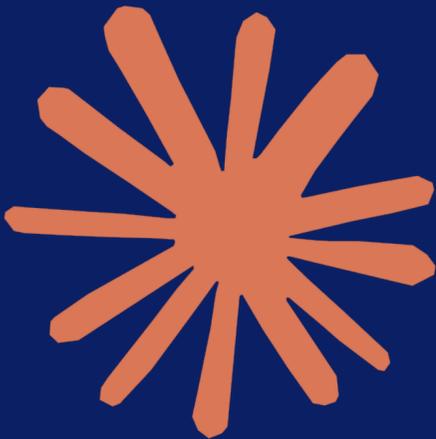
```
## How Skills Work
```

1. Brian asks to do something
2. Claude checks this index for matching skill
3. If found, Claude loads that skill file and follows procedure
4. If not found, Claude figures it out and offers to create a new skill

```
## Available Skills
```

Skill	Trigger Phrases	File
Process Inbox	"process inbox", "what's in inbox"	[[skills/process-
Write Blog Post	"write a blog post", "draft post about"	[[skills/w
Create LinkedIn Post	"linkedin post", "post on linkedin"	[[skills/
Write Abstract	"write abstract", "conference abstract"	[[skills/wr
Process Meeting	"process this transcript", "meeting notes from"	[[
Add to Website	"add to website", "update my website", "I got quoted	[[skills/add-to-website]]

AI a



Chat Cowork **Code**

Write weekly blog post ▼ ▶ Preview ▼

Let's work on this week's blog post.

Let me load up the relevant context.

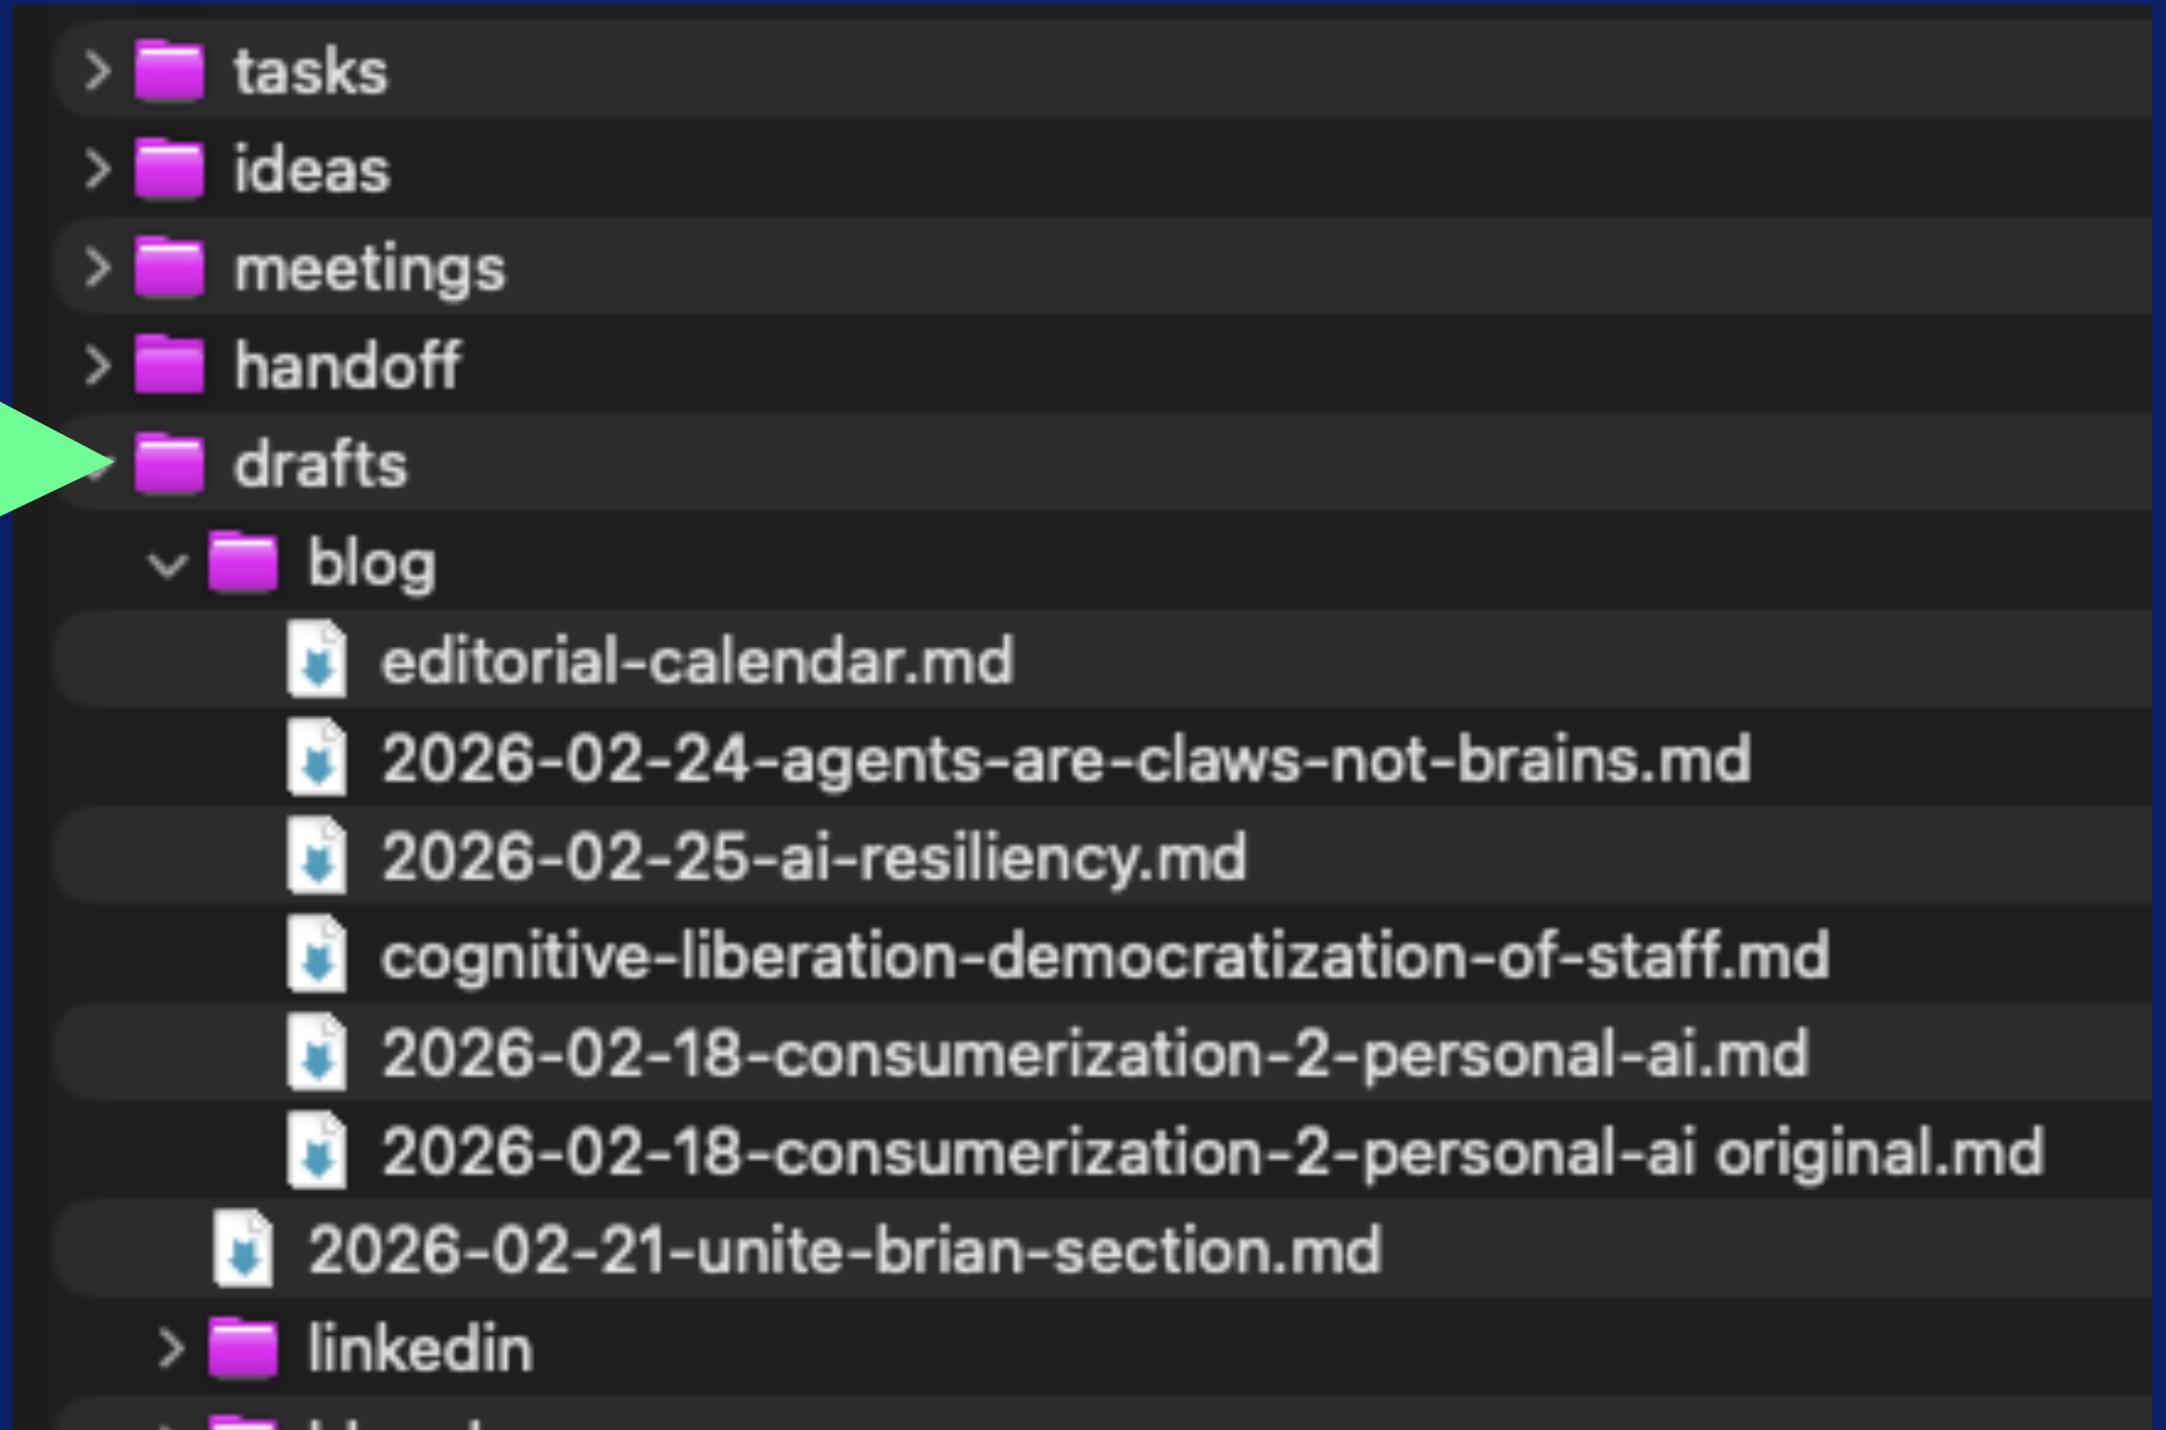
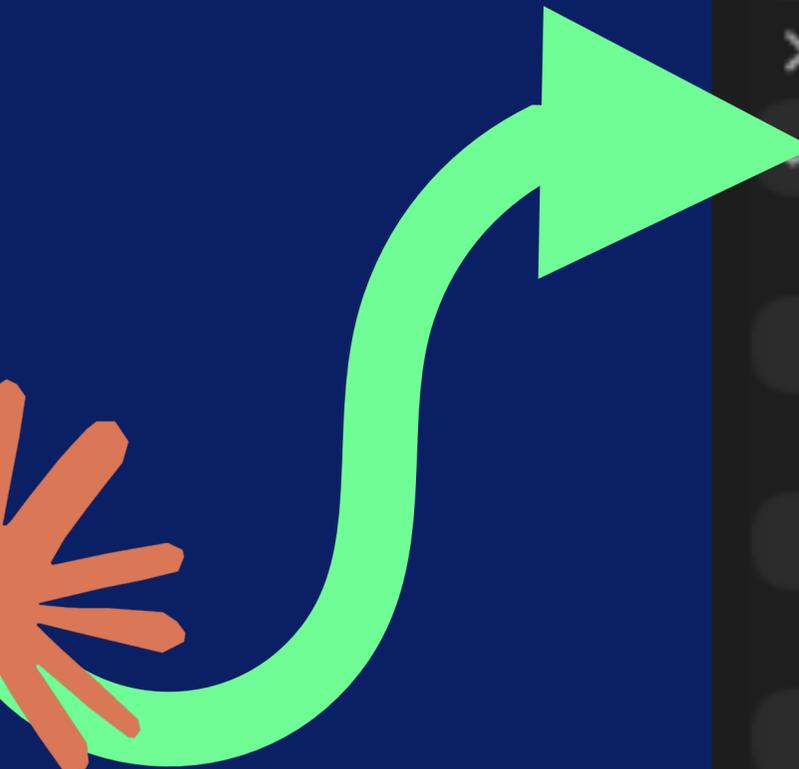
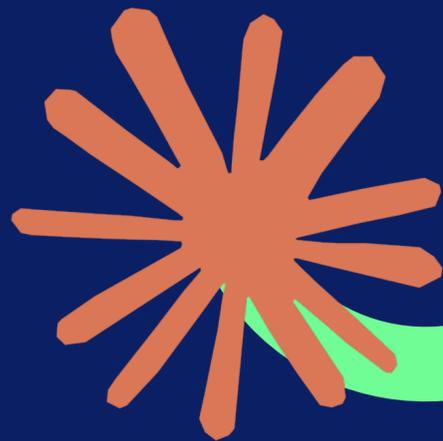
```
▼ Read three markdown documents about AI stack economics and thinking notes
- Read /Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/me/thinking.md
- Read /Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/tasks/tasks.md
  Read 135 lines
- Glob drafts/**/*.md
  /Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/drafts/2026-02-11-deep-research-prompt-token-economics.md
  /Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/drafts/blog/2026-02-18-consumerization-2-personal-ai original.md
  /Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/drafts/cto-role-definition.md
- Show 5 more
```

Here's where things stand:

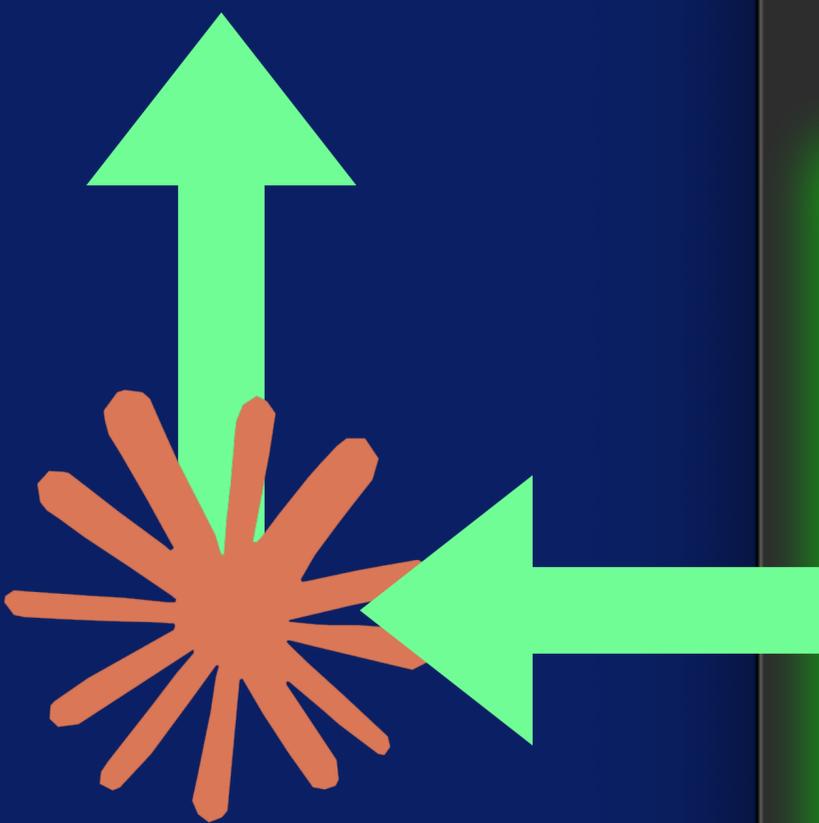
Last published: "Skills are all you need" (March 12)

Overdue in the editorial calendar:

AI as my cognitive extension



AI a



Read `/Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/me/thinking.md`

Read `/Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/tasks/tasks.md`

Read 135 lines

Glob `drafts/**/*.md`

`/Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/drafts/2026-02-11-deep-research-prompt-token-economics.md`

`/Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/drafts/blog/2026-02-18-consumerization-2-personal-ai-original.md`

`/Users/brian/Library/Mobile Documents/com~apple~CloudDocs/bmad/drafts/cto-role-definition.md`

Show 5 more

Here's where things stand:

Last published: "Skills are all you need" (March 12)

Overdue in the editorial calendar:

1. "The bitter lesson, revisited" -- framework complete, target was week of March 3
2. "The five levels of AI in knowledge work" -- framework complete, target was week of March 10
3. "Enterprise invariants" -- idea complete, target was March 10-17 window

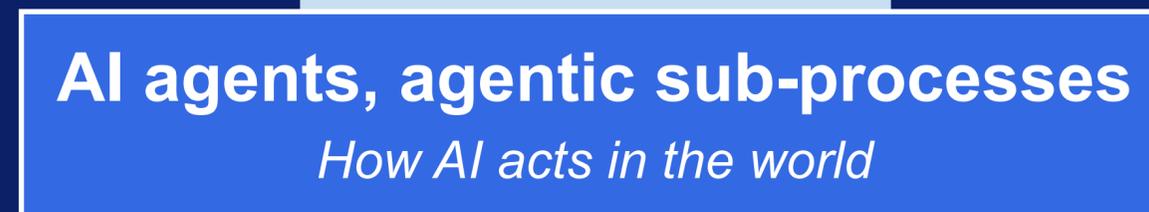
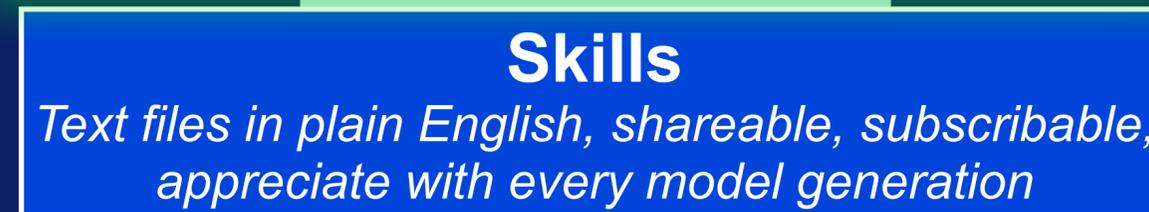
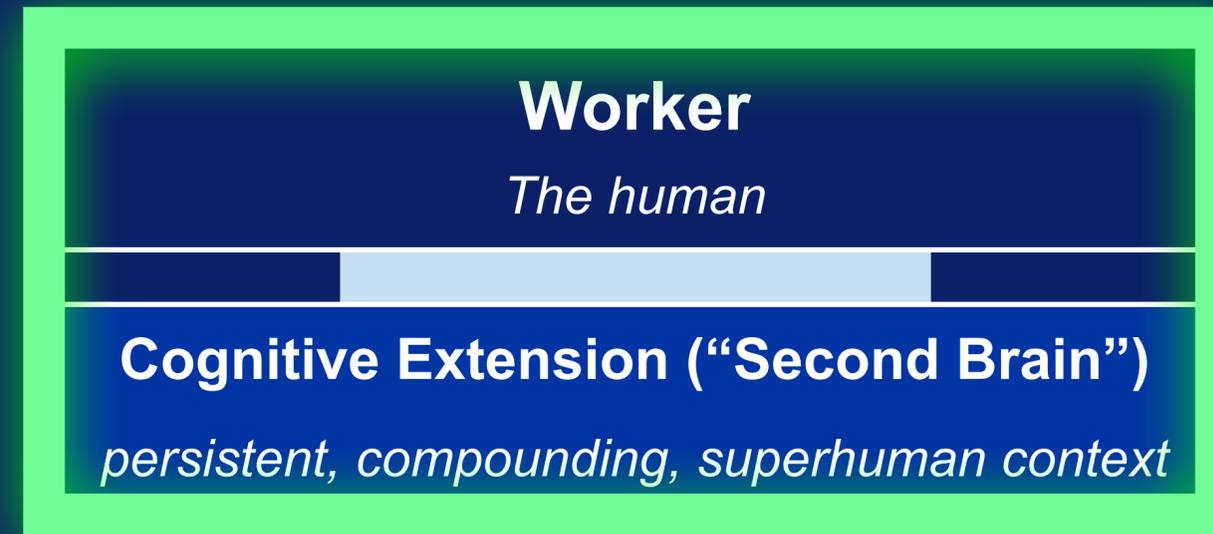
Fresh seeds from today: AI stack value layers, FDE companies as consulting firms

Which post do you want to work on? The bitter lesson and five levels are both framework-complete and just need drafting. Enterprise invariants is also ready. Or if something else is pulling at you right now.

The cognitive stack of AI-enhanced knowledge work

“80%” Invisible
Thinking

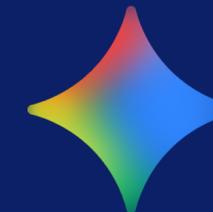
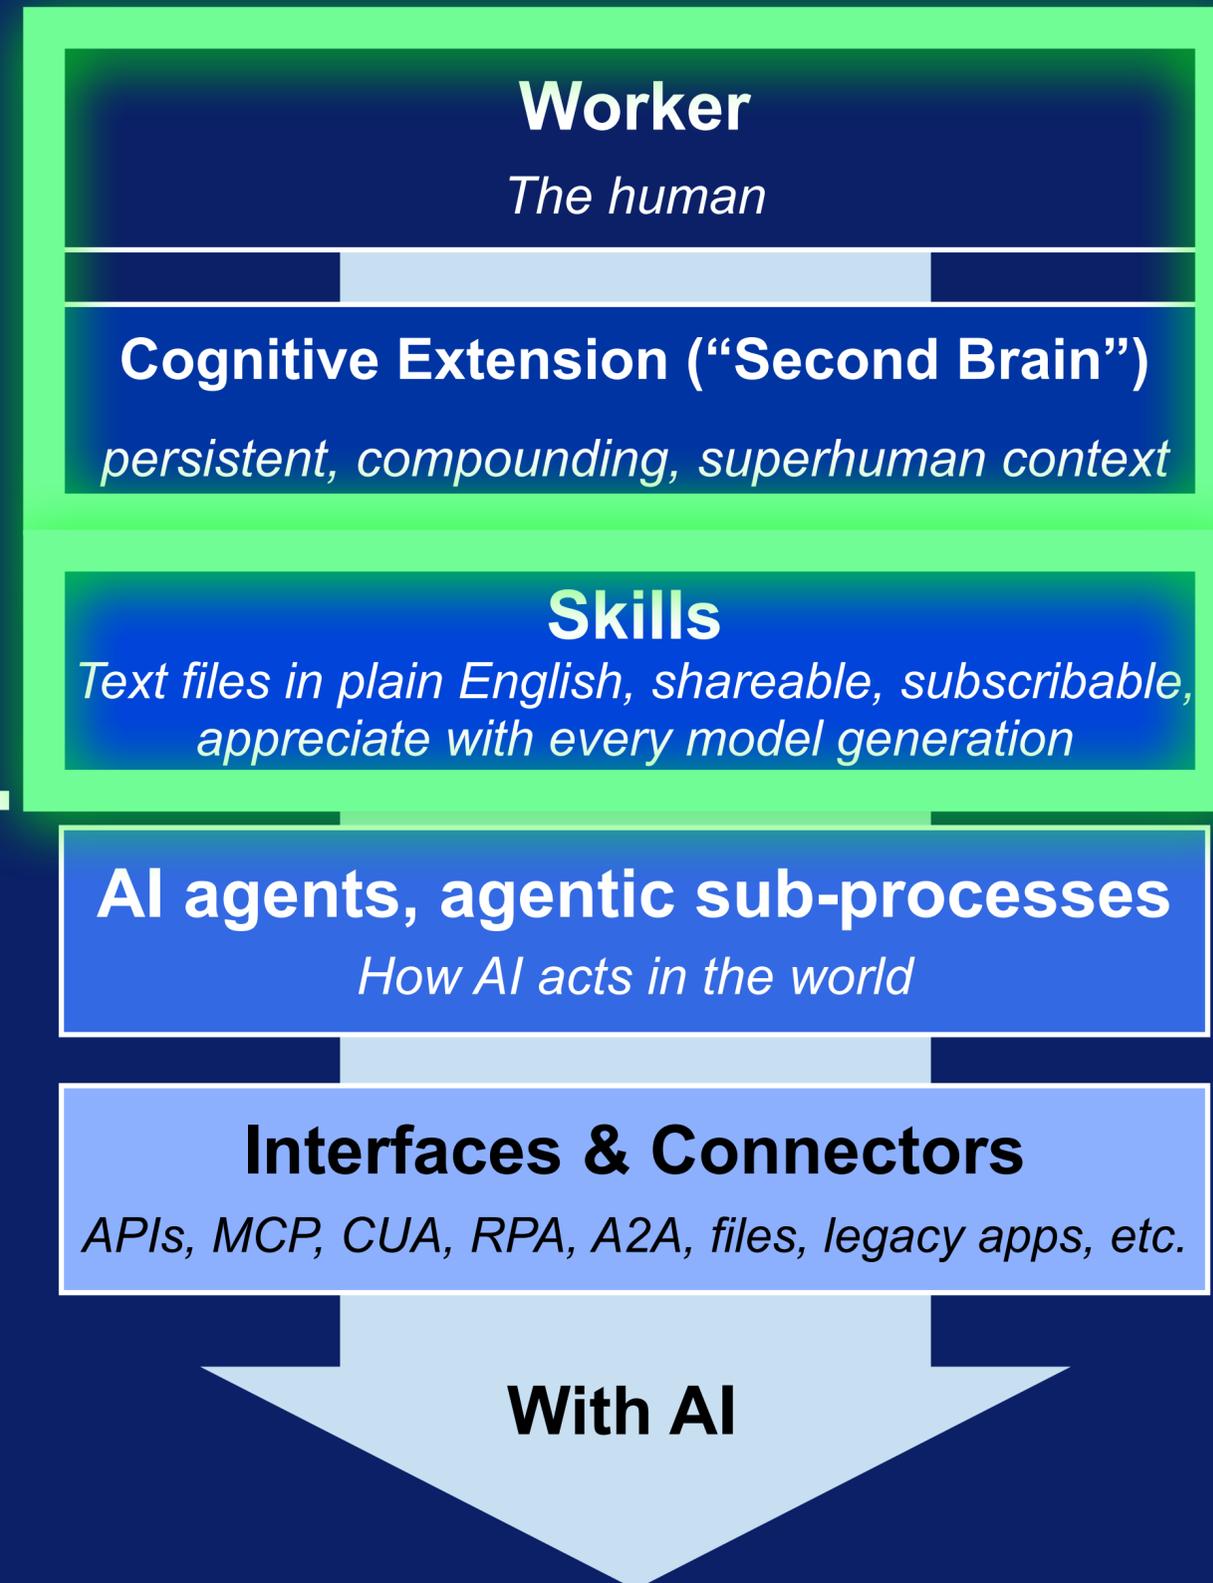
Doing
“20%” Visible



The cognitive stack of AI-enhanced knowledge work

“80%” Invisible
Thinking

Doing
“20%” Visible



How does the AI know what to do? Skills!

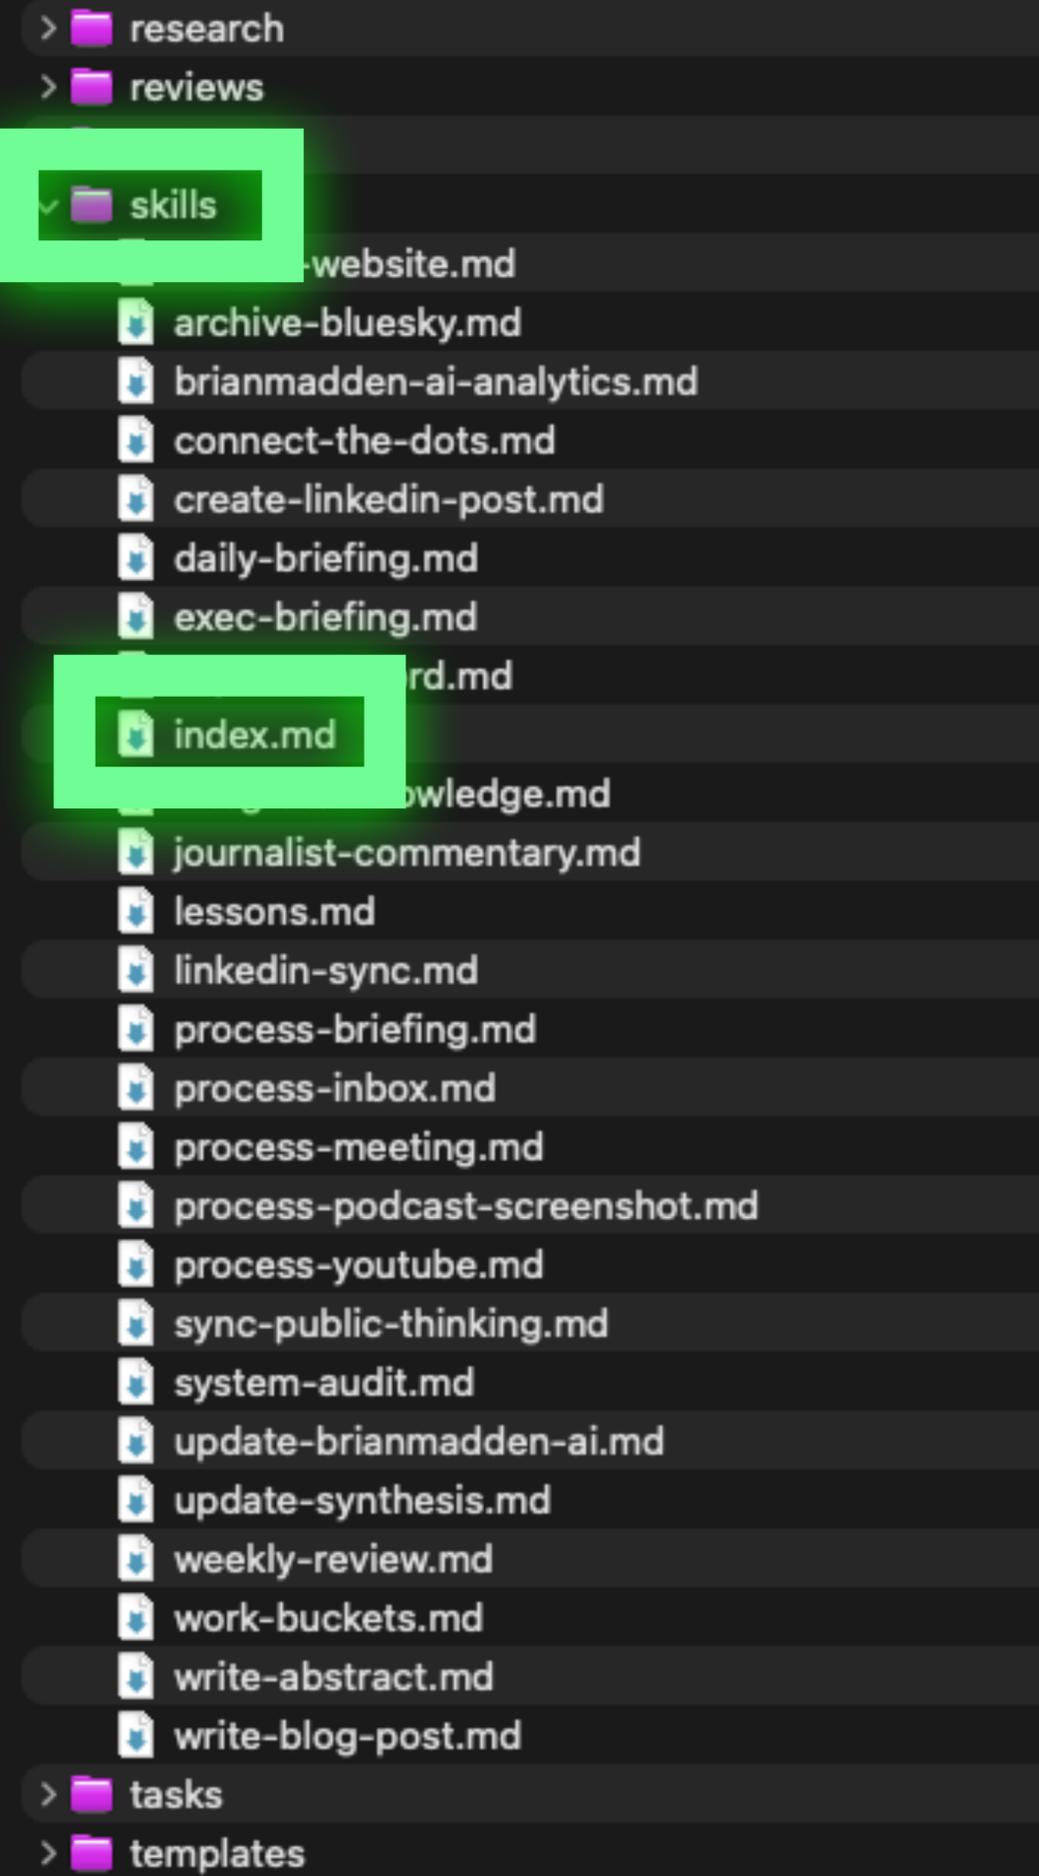
“Skills” are text files that tell the AI how to do something

They connect the “brain” to the “claws” (thinking to doing)

NOT engineering. If you can write instructions for a colleague, that’s a skill!

(Really the AI mostly writes its own skills)

Better AI models = better execution of existing skills!



Skills Index

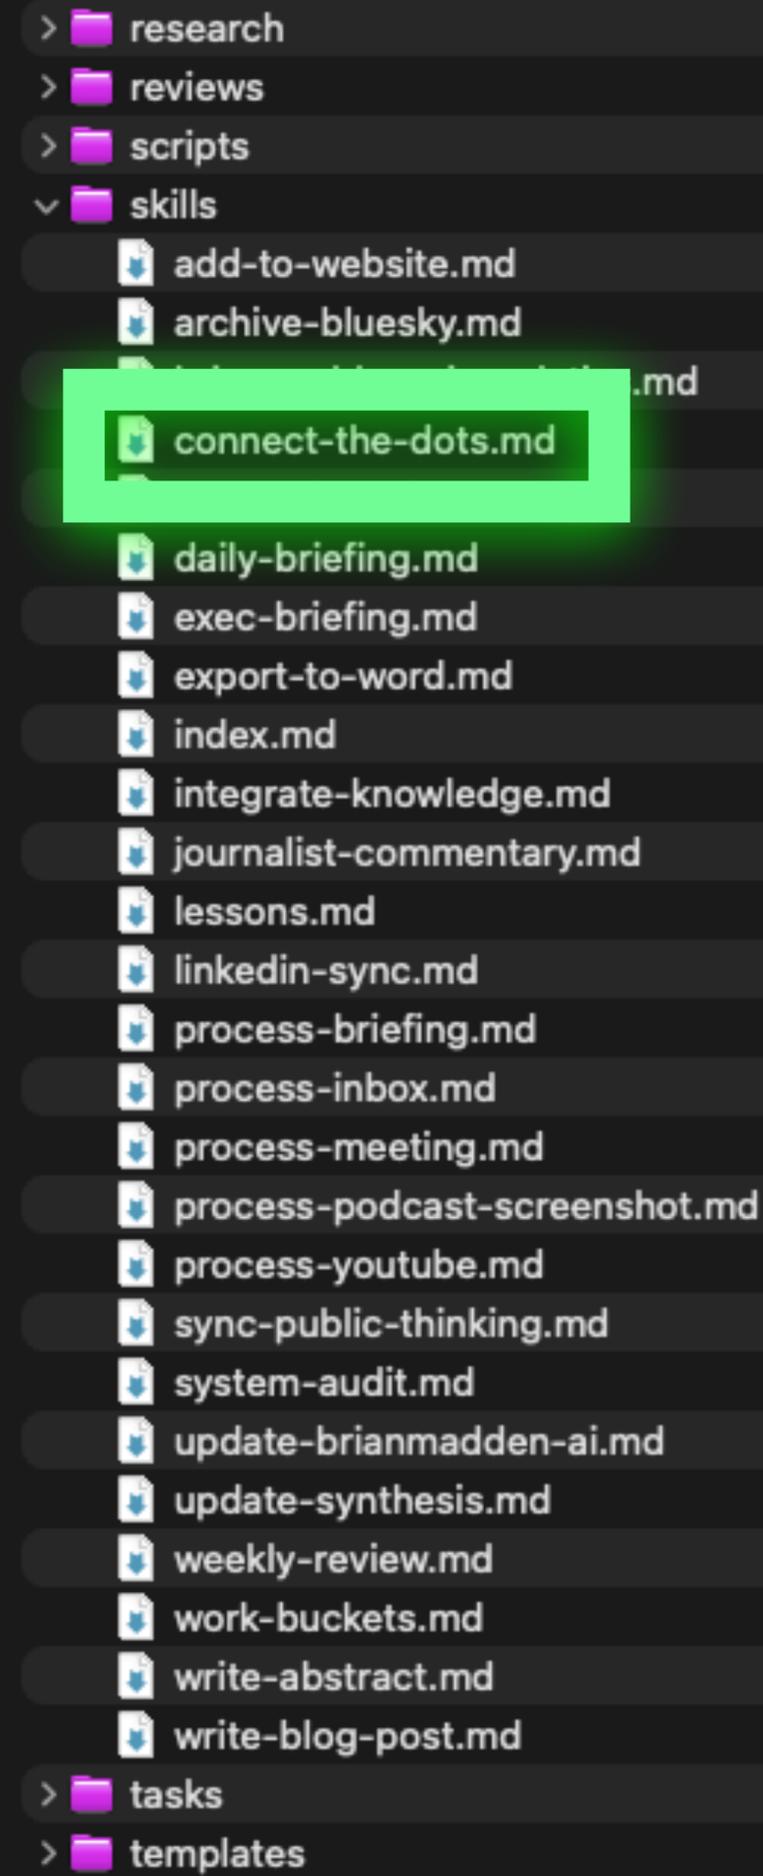
Skills are procedures Claude can follow. When Brian asks to do something, ch

How Skills Work

1. Brian asks to do something
2. Claude checks this index for matching skill
3. If found, Claude loads that skill file and follows procedure
4. If not found, Claude figures it out and offers to create a new skill

Available Skills

Skill	Trigger Phrases	File
Process Inbox	"process inbox", "what's in inbox"	[[skills/process-inbo
Write Blog Post	"write a blog post", "draft post about"	[[skills/write
Create LinkedIn Post	"linkedin post", "post on linkedin"	[[skills/crea
Write Abstract	"write abstract", "conference abstract"	[[skills/write-
Process Meeting	"process this transcript", "meeting notes from"	[[skil
Add to Website	"add to website", "update my website", "I got quoted in",	
Journalist Commentary	"journalist questions", "press inquiry", "PR reque	
Update Synthesis	"new blog post added", "update synthesis", "I published	
Daily Briefing	"daily briefing", "what should I read today", "process fe	
Process Briefing	"process briefing", "review briefing", "catch me up", "	
Weekly Review	"weekly review", "time to dream", "consolidate the brain",	
Exec Briefing	"write an exec note", "Nancy should know", "drop a note to	
Process YouTube	"process youtube", "youtube transcript", YouTube URL in	
Process Podcast Screenshot	podcast screenshot in inbox, "process podcast	



Deep knowledge synthesis across the entire BMAD system. This skill reads everything, finds missing connections, hunts for gaps, searches the web for adjacent ideas and counterarguments, and proposes new intellectual territory. The system audit keeps the house clean. This skill pushes the walls out.

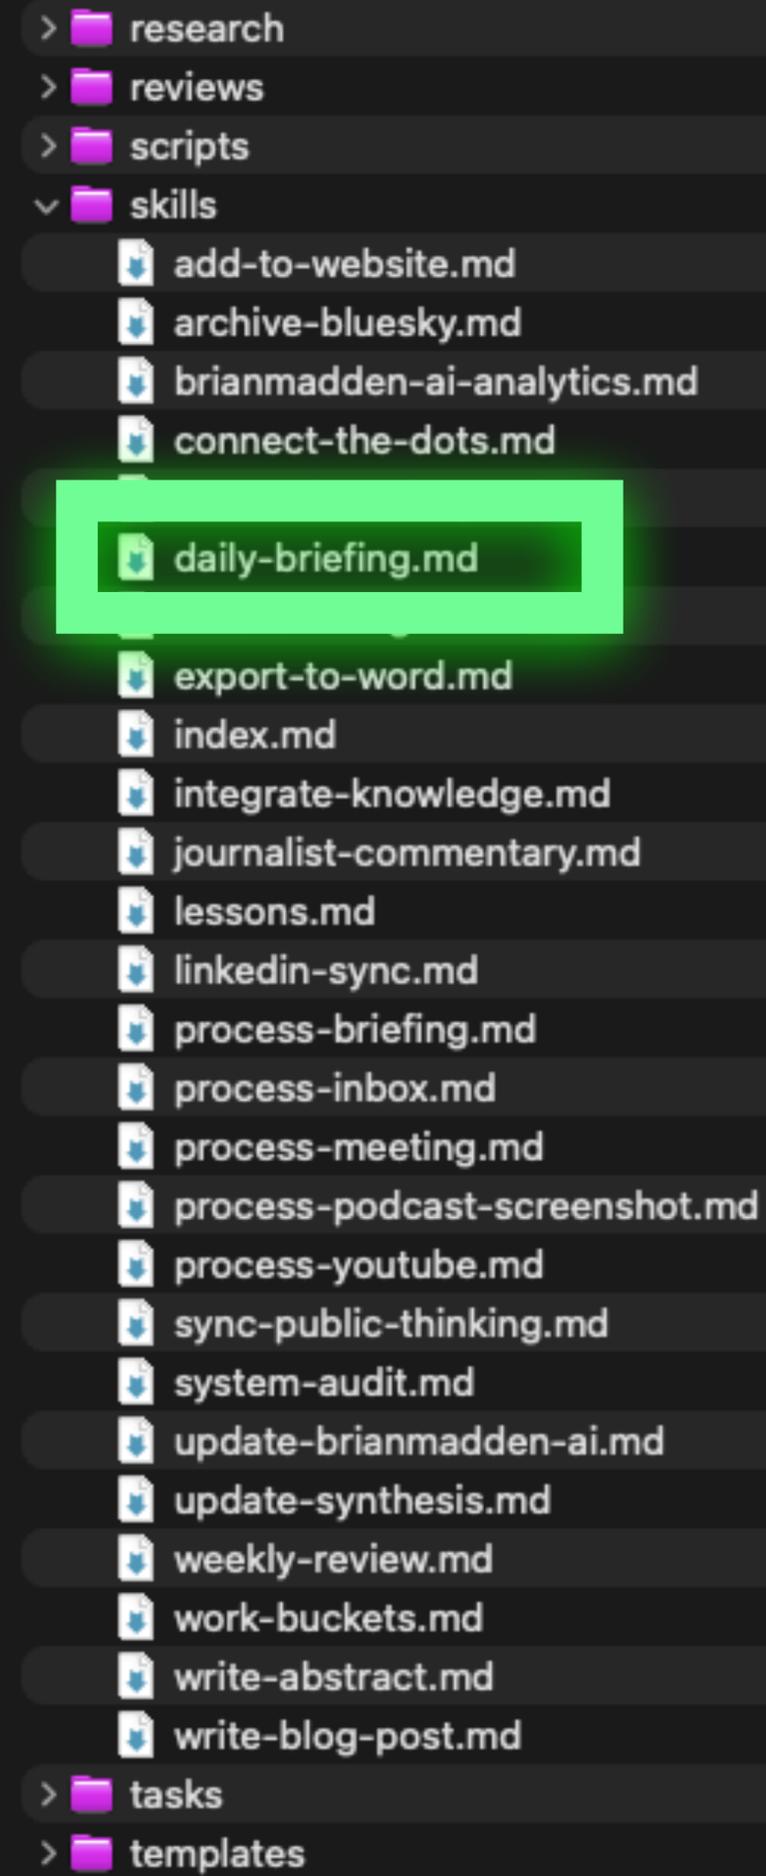
Triggers

- "connect the dots"
- "what am I missing"
- "find the gaps"
- "push the thinking"
- "what should I be reading"
- "challenge my positions"
- "where are the blind spots"
- Recommended cadence: weekly, or whenever Brian feels intellectually stuck

Purpose

A second brain that only stores and organizes is just a filing cabinet. This skill makes it **think**. It reads across every level of the knowledge hierarchy, finds connections that haven't been made, identifies arguments that need strengthening, discovers external voices worth engaging with, and surfaces the questions Brian should be asking but isn't.

This is the skill that compounds the system's intellectual value over time.



Generates a personalized daily briefing from RSS feeds, personalized to Brian's context. Runs automatically via GitHub Actions at 6 AM CET, Monday through Friday. Can also be triggered manually.

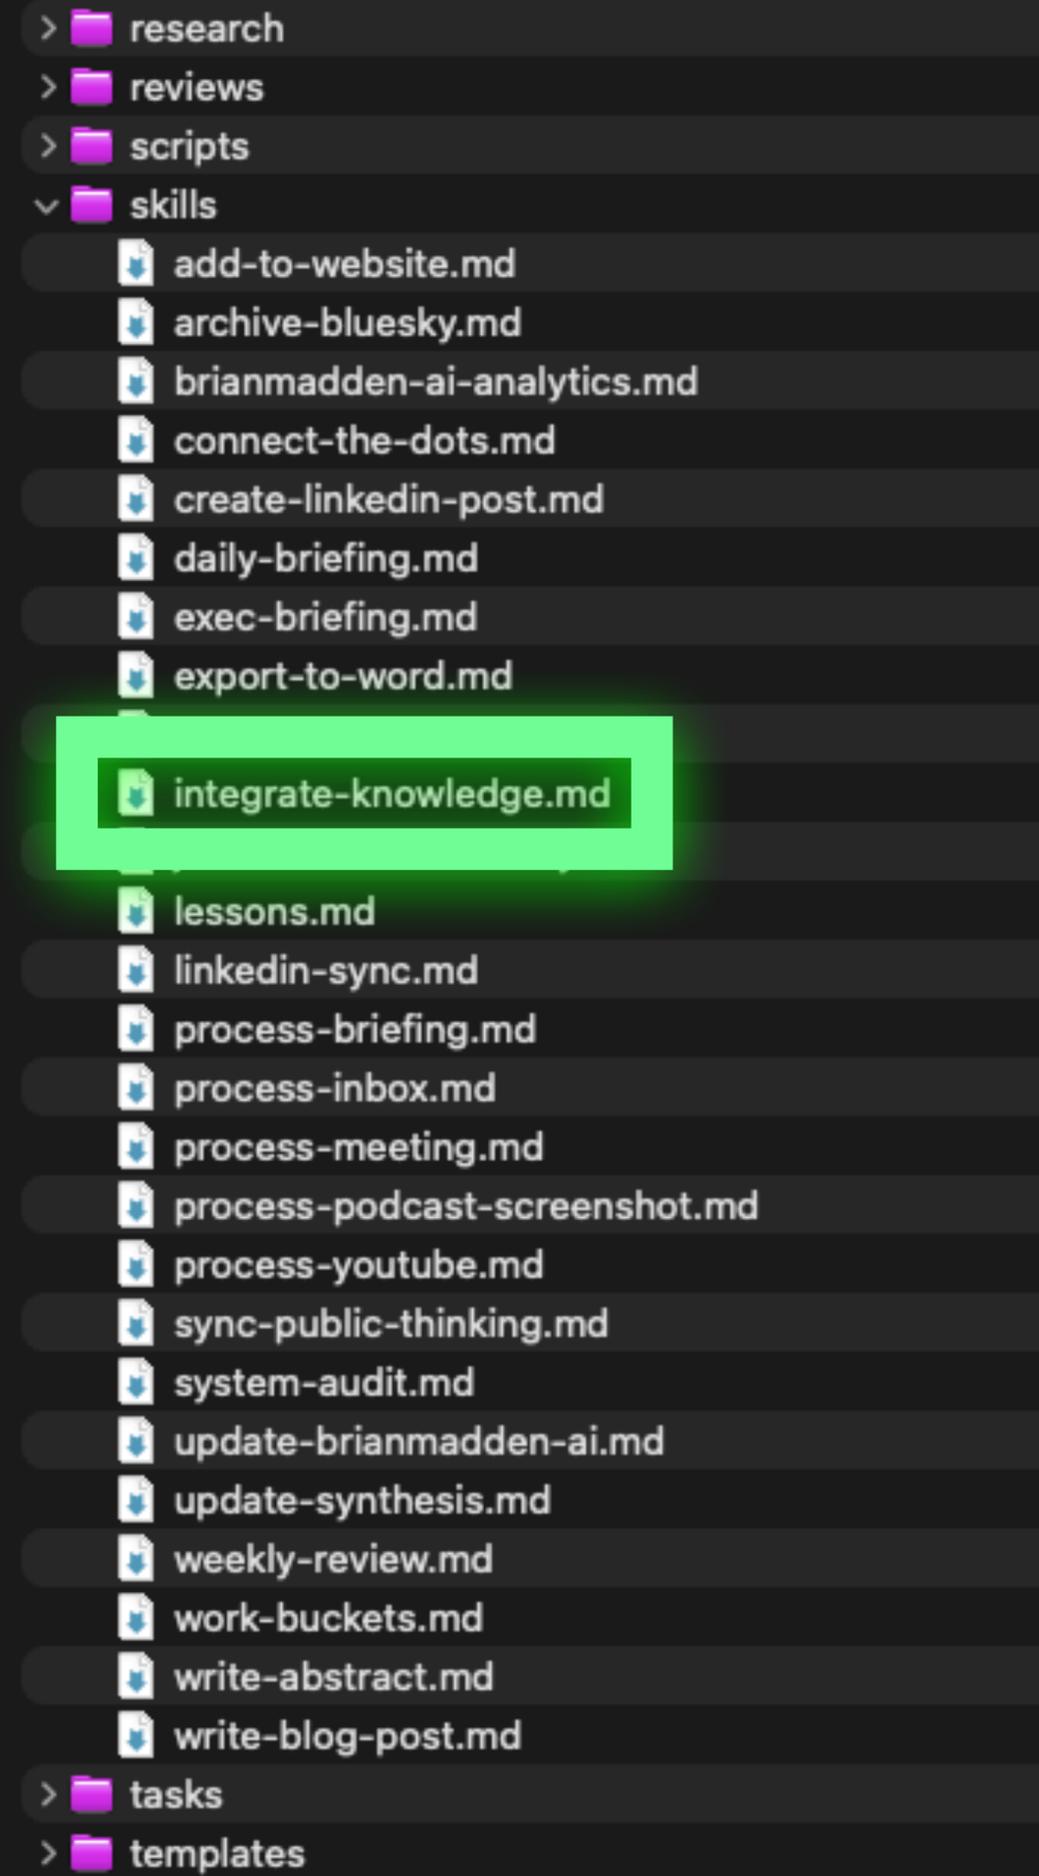
Automated flow (GitHub Actions)

The workflow in ``.github/workflows/daily-briefing.yml``:

1. Runs at 05:00 UTC (06:00 CET) weekdays
2. **Fetches emails** from iCloud IMAP "For AI Processing" folder via ``.scripts/fetch-emails.py`` (optional-continues if no credentials or no emails)
3. **Skips if today's briefing already exists**—won't overwrite
4. Fetches all RSS feeds from ``.feeds/sources.yml``
5. For high-priority articles, extracts full content using reader mode
6. Loads any emails from ``.briefings/email-intake/`` (placed by step 2)
7. Loads BMAD context:
 - ``.me/thinking.md``—Brian's current thinking
 - ``.work/arcs/``—Active narrative threads
 - ``.work/citrix/strategy.md``—Current job context
 - ``.work/intel/index.md``—Key company tracking
8. Calls Claude API with articles + emails + context
9. Writes output to ``.briefings/YYYY-MM-DD.md``
10. Cleans up email intake folder
11. Commits and pushes

Standalone briefings

Each day's briefing is independent. Prior briefings are left in



Filing is organized forgetting. This skill ensures new material doesn't just into active thinking. Every new source should update the files where Brian's

Inputs

- One or more source files (research outputs, transcripts, strategy docs, de
- Already placed in their storage location (or being placed now)

Procedure

1. Read the source material

Understand what's in it. Identify the 3-10 most important insights, facts, f

2. Map connections to active thinking

For each significant insight, check which active files it connects to:

Active file	What to look for
`me/synthesis.md`	Does it reinforce, extend, or challenge a published fr
`me/thinking.md`	Does it inform an evolving position? Add new evidence?
`work/arcs/*`	Does it feed a narrative arc? Add supporting evidence?
`work/citrix/strategy.md`	Does it change competitive positioning or add
`work/intel/*`	Does it update what we know about a company or market?
`ideas/*`	Does it connect to or strengthen a seed idea?
`tasks/tasks.md`	Does it create new action items or change priorities?

3. Propose updates

```
> research
> reviews
> scripts
skills
  add-to-website.md
  archive-bluesky.md
  brianmadden-ai-analytics.md
  connect-the-dots.md
  create-linkedin-post.md
  daily-briefing.md
  exec-briefing.md
  export-to-word.md
  index.md
  integrate-knowledge.md
  journalist-commentary.md
  lessons.md
  linkedin-sync.md
  process-briefing.md
  process-inbox.md
  process-podcast-screenshot.md
  sync-public-thinking.md
  system-audit.md
  update-brianmadden-ai.md
  update-synthesis.md
  weekly-review.md
  work-buckets.md
  write-abstract.md
  write-blog-post.md
```

Inputs

- A screenshot (PNG/JPG) from a phone showing a podcast playing
- The screenshot contains: podcast name, episode title/number, playback position, or progress bar)

Procedure

1. Read the screenshot

Extract from the image:

- **Podcast name** (e.g., "The Artificial Intelligence Show")
- **Episode number and/or title** (e.g., "#196: SaaSocalypse")
- **Playback timestamp**—elapsed time, remaining time, or estimate from progress bar
- Calculate total episode duration if both elapsed and remaining are visible

2. Identify the podcast in the system

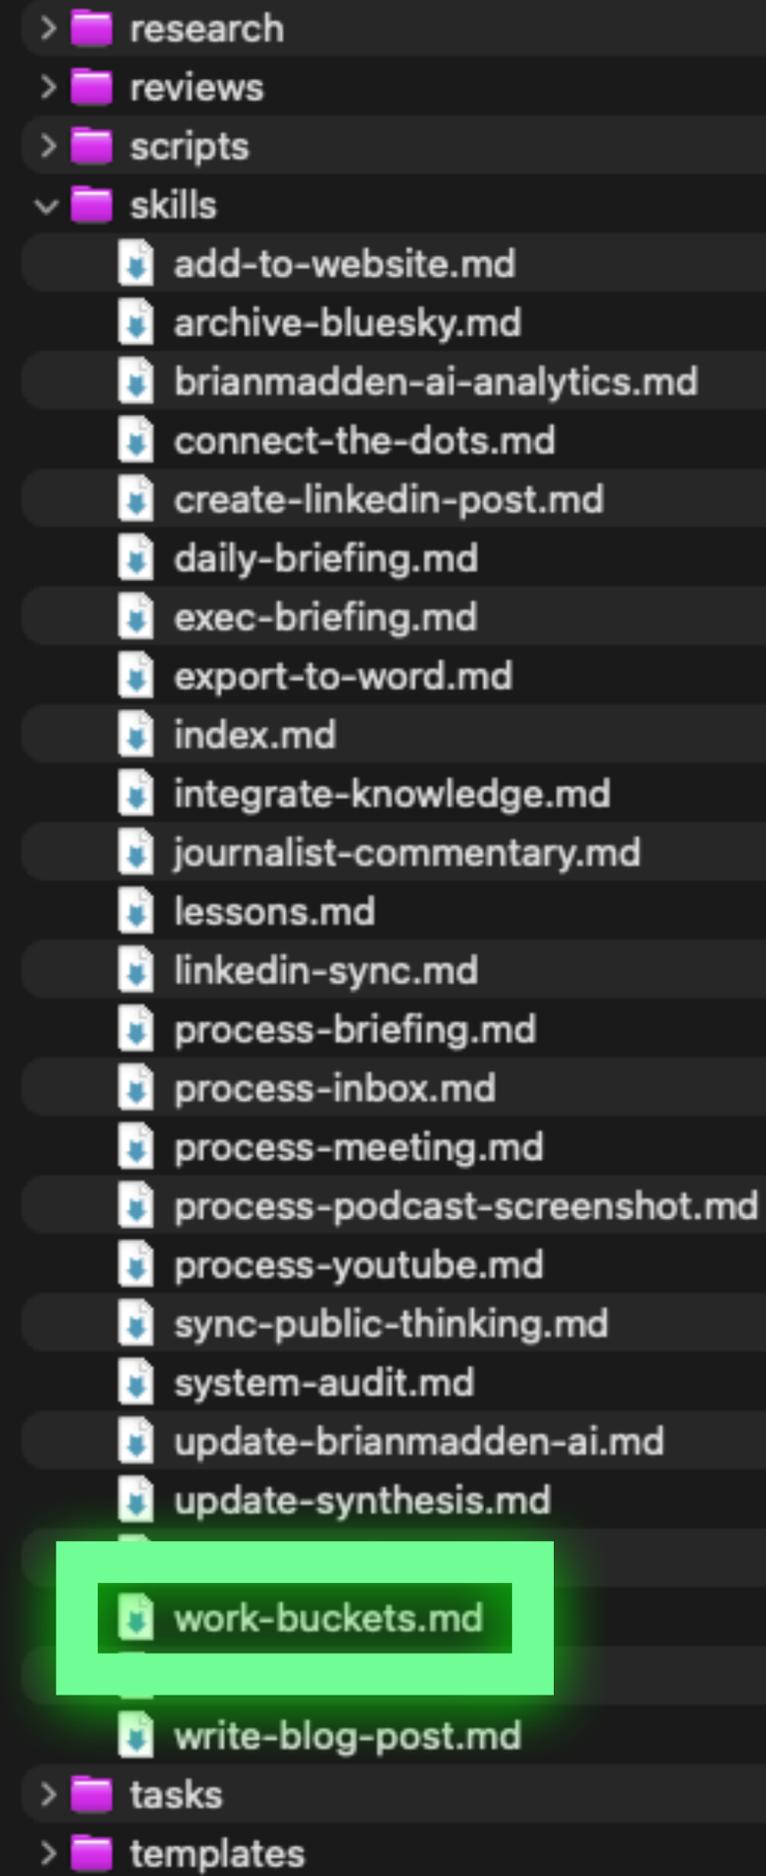
Check `feeds/sources.yaml` under `podcasts_with_transcripts`, `podcasts_via` sections. If the podcast is known, note:

- Any transcript source URL pattern
- Priority level
- Whether transcripts are published on the podcast's website

If the podcast is not in the system, note that and proceed—but offer to add afterward.

3. Find the episode and transcript

If the podcast publishes transcripts on their website (noted in `feeds/s



What Buckets Are

Lightweight collector files in `work/buckets/`. Each bucket is a markdown file that links to relevant materials with brief context on *why* each item matters for an event. Items can appear in multiple buckets.

Buckets are ephemeral—created when something is on the horizon, materials are collected, and then when it's time to prep, everything gets processed. After the event, archive the bucket.

Procedure

Adding to a bucket

1. If bucket file doesn't exist, create it in `work/buckets/` with slug name
2. Add the wiki-link with a brief context annotation (why this item matters)
3. Items can go in multiple buckets—just add the link to each

Prepping from a bucket

When Brian says "prep for [X]":

1. Load the bucket file
2. Load all linked documents
3. Synthesize into talking points, prep notes, or whatever format the event needs
4. Output to the appropriate location (could be the bucket's Notes section, or a separate file)

Bucket lifecycle

- **Create** when an event, meeting, initiative, or deadline is on the horizon
- **Accumulate** materials as they come in over days/weeks

anthropics / skills

Code Issues 155 Pull requests 349 Discussions More

Files main skills / skills

zack-anthropic skill-creator: drop ANTHR... b0cbd3d · 2 weeks ago

Name	Last commit message	Last commit date
..		
algorithmic-art	Move example skills into d...	3 months ago
brand-guidelines	Move example skills into d...	3 months ago
canvas-design	Move example skills into d...	3 months ago
claude-api	Add claude-api skill (#515)	2 weeks ago
doc-coauthoring	Add doc-coauthoring skill...	3 months ago
docx	chore: export latest skills ...	3 weeks ago
frontend-design	Add doc-coauthoring skill...	3 months ago
internal-comms	Move example skills into d...	3 months ago
mcp-builder	Move example skills into d...	3 months ago
pdf	Update skill-creator and ...	last month
pptx	Update skill-creator and ...	last month
skill-creator	skill-creator: drop ANTHR...	2 weeks ago
slack-gif-creator	Add doc-coauthoring skill...	3 months ago
theme-factory	Move example skills into d...	3 months ago
web-artifacts-builder	Move example skills into d...	3 months ago
webapp-testing	Move example skills into d...	3 months ago
xlsx	Update skill-creator and ...	last month

VoltAgent / awesome-agent-skills

Code Issues Pull requests 15 Actions Projects More

Watch 98

Claude Code Skills and 500+ agent skills from official dev teams and the community, compatible with Codex, Antigravity, Gemini CLI, Cursor and others.

github.com/VoltAgent/voltagent

MIT license

Contributing

11.6k stars 1.1k forks 98 watching 21 Branches

0 Tags Activity Custom properties

Public repository

main

Go to file + <> Code

necatiozmen Merge pull request #221 fro... ce4dc1d · 14 hours ago

CONTRIBUTING.md	update: contributing guide	2 months ago
LICENSE	Initial commit	5 months ago
README.md	Merge pull request #221 f...	14 hours ago

README Contributing MIT license

Opencode

Gemini

Claude Code

Codex

Antigravity

Agent Skills Collection
by official dev teams & community

Hand-picked, not bulk-imported

voltagent

openai / skills

Code Issues 36 Pull requests 109 Actions Security Insights

Files main skills / skills / .curated

fjord-oai Update playwright-interactive exam... ecc1fb8 · last week

Name	Last commit message	Last commit date
..		
aspnet-core	Add aspnet and winui cur...	2 weeks ago
chatgpt-apps	remove image file require...	last week
cloudflare-deploy	Add prerequisites and tro...	last month
develop-web-game	Add curated default prom...	last month
doc	Add curated default prom...	last month
figma-implement-design	Add curated default prom...	last month
figma	Add curated default prom...	last month
gh-address-comments	Add curated default prom...	last month
gh-fix-ci	Add curated default prom...	last month
imagegen	Add curated default prom...	last month
jupyter-notebook	Add curated default prom...	last month
linear	Add curated default prom...	last month
netlify-deploy	Add prerequisites and tro...	last month
notion-knowledge-capt...	Add curated default prom...	last month
notion-meeting-intellig...	Add curated default prom...	last month
notion-research-docu...	Add curated default prom...	last month
notion-spec-to-implem...	Add curated default prom...	last month
openai-docs	Update curated OpenAI D...	2 weeks ago
pdf	Add curated default prom...	last month
playwright-interactive	Update playwright-interac...	last week

sickn33 / antigravity

Code Issues Pull requests

Watch 225

The Ultimate Collection of 1 Code/Antigravity/Cursor. Ba agents including official skill

github.com/sickn33/antigravity

MIT license

Contributing

Security policy

25.2k stars 4.3k forks

83 Tags Activity

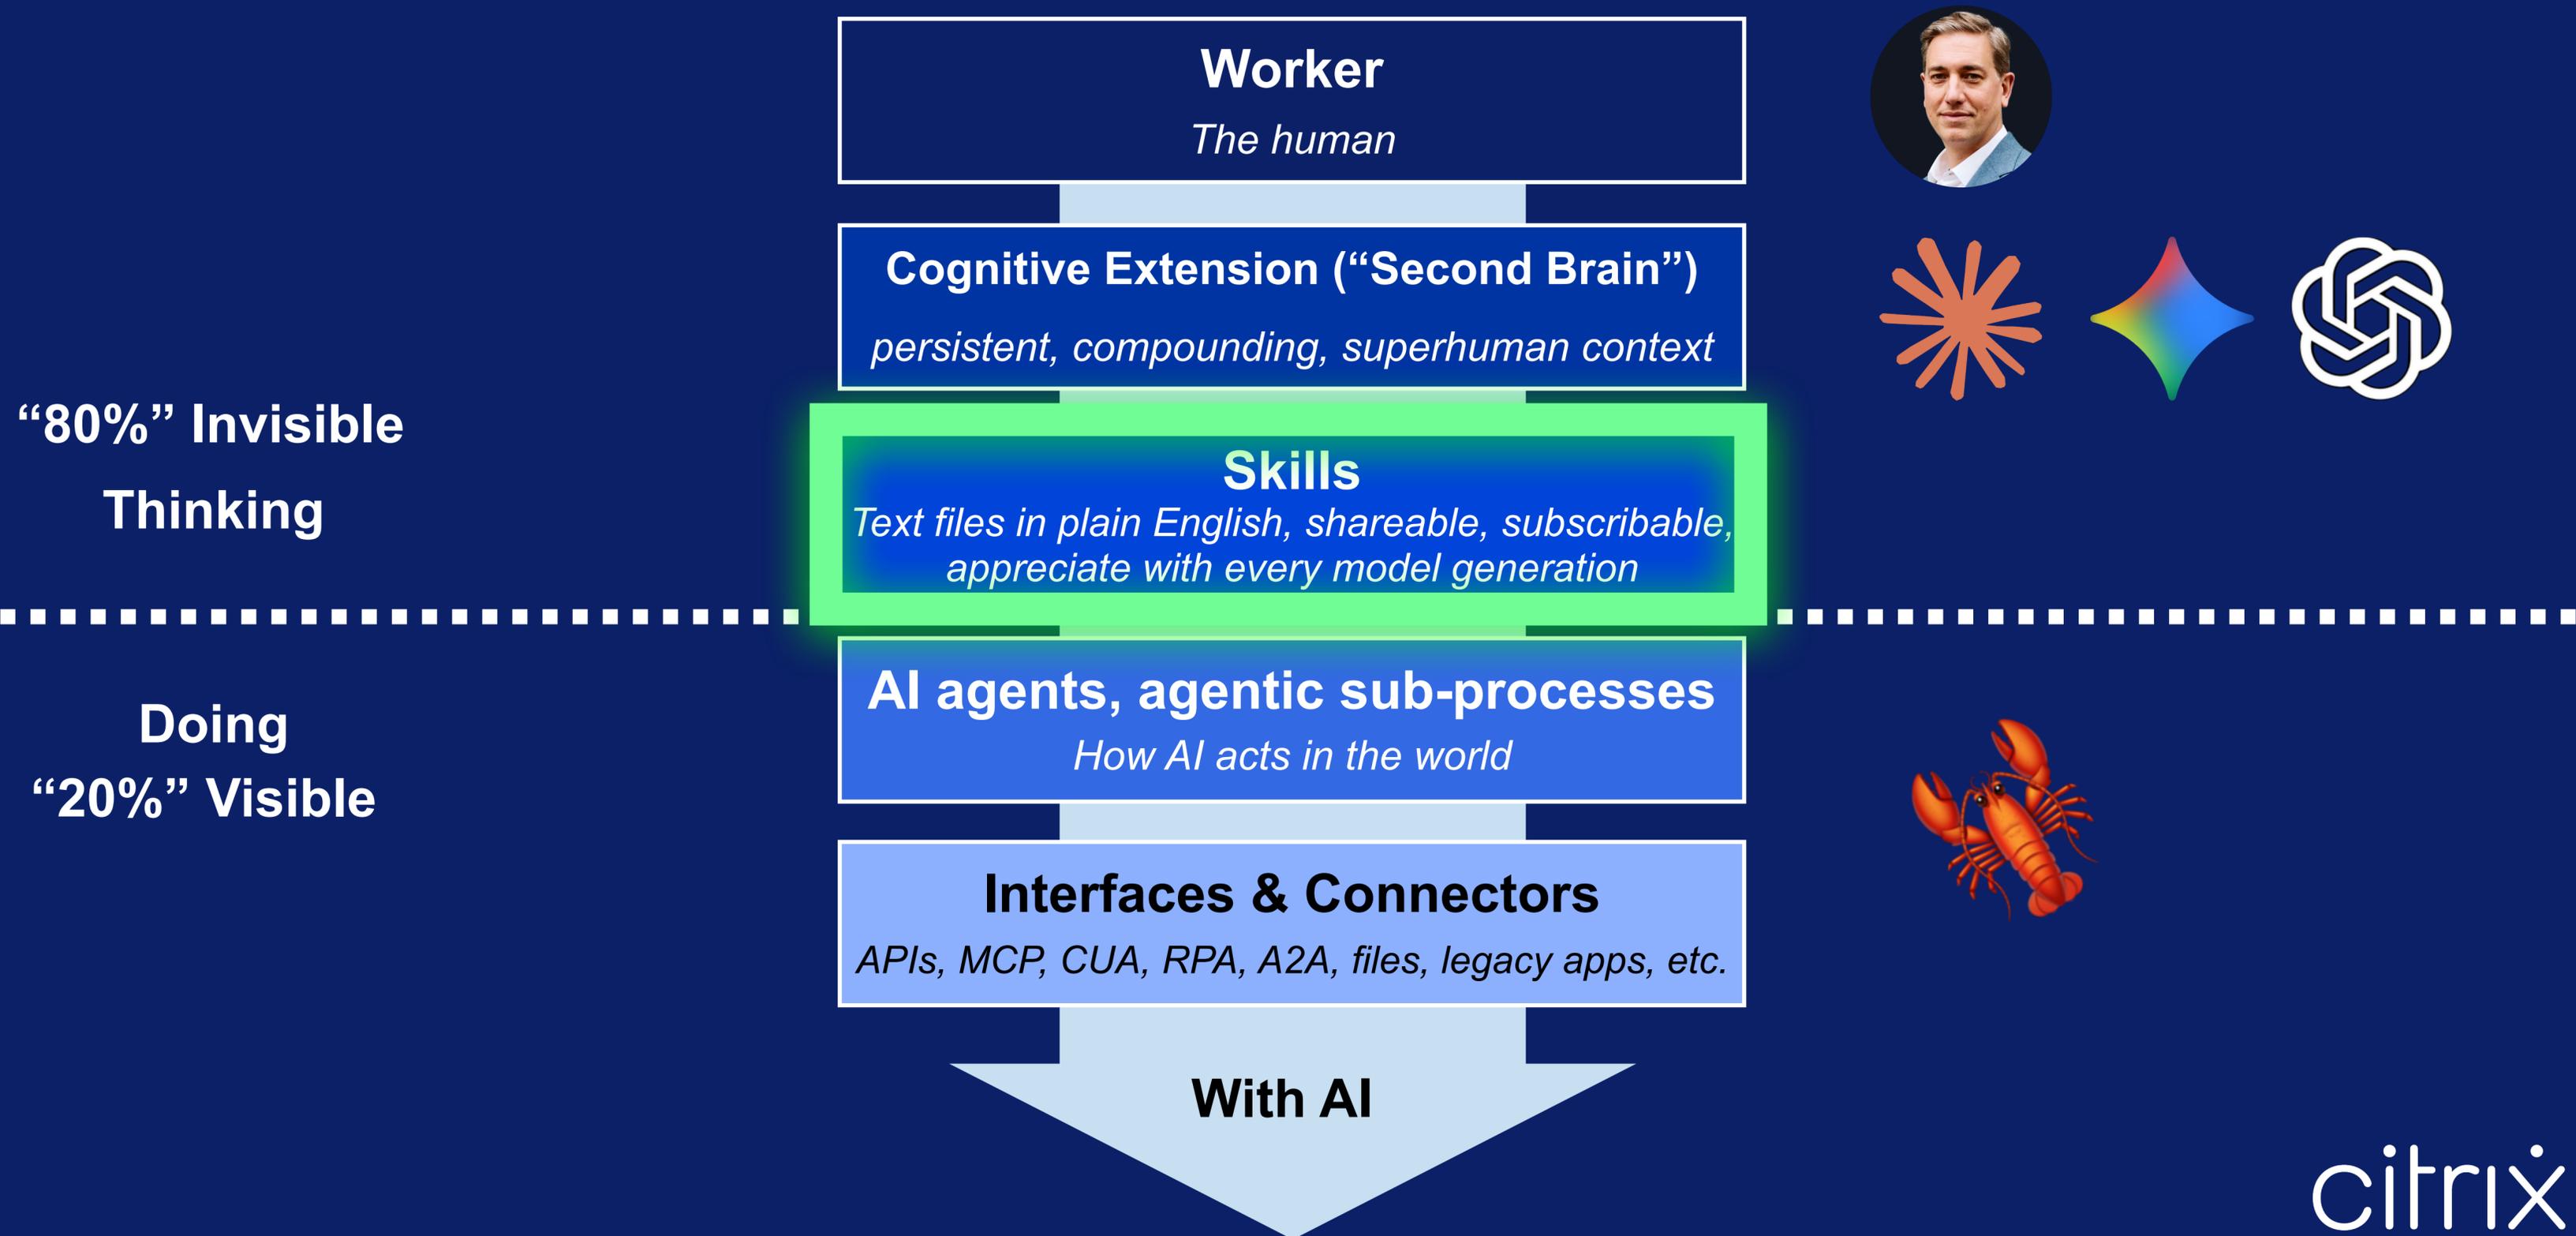
Public repository

main

sickn33 chore: release v8

- .claude-plugin
- .github
- apps/web-app
- assets
- data
- docs
- docs_zh-CN
- scripts
- skill_categorization
- skills

The cognitive stack of AI-enhanced knowledge work



automations

automations

automations

automations

automations

This is **STILL** not about automations

automations

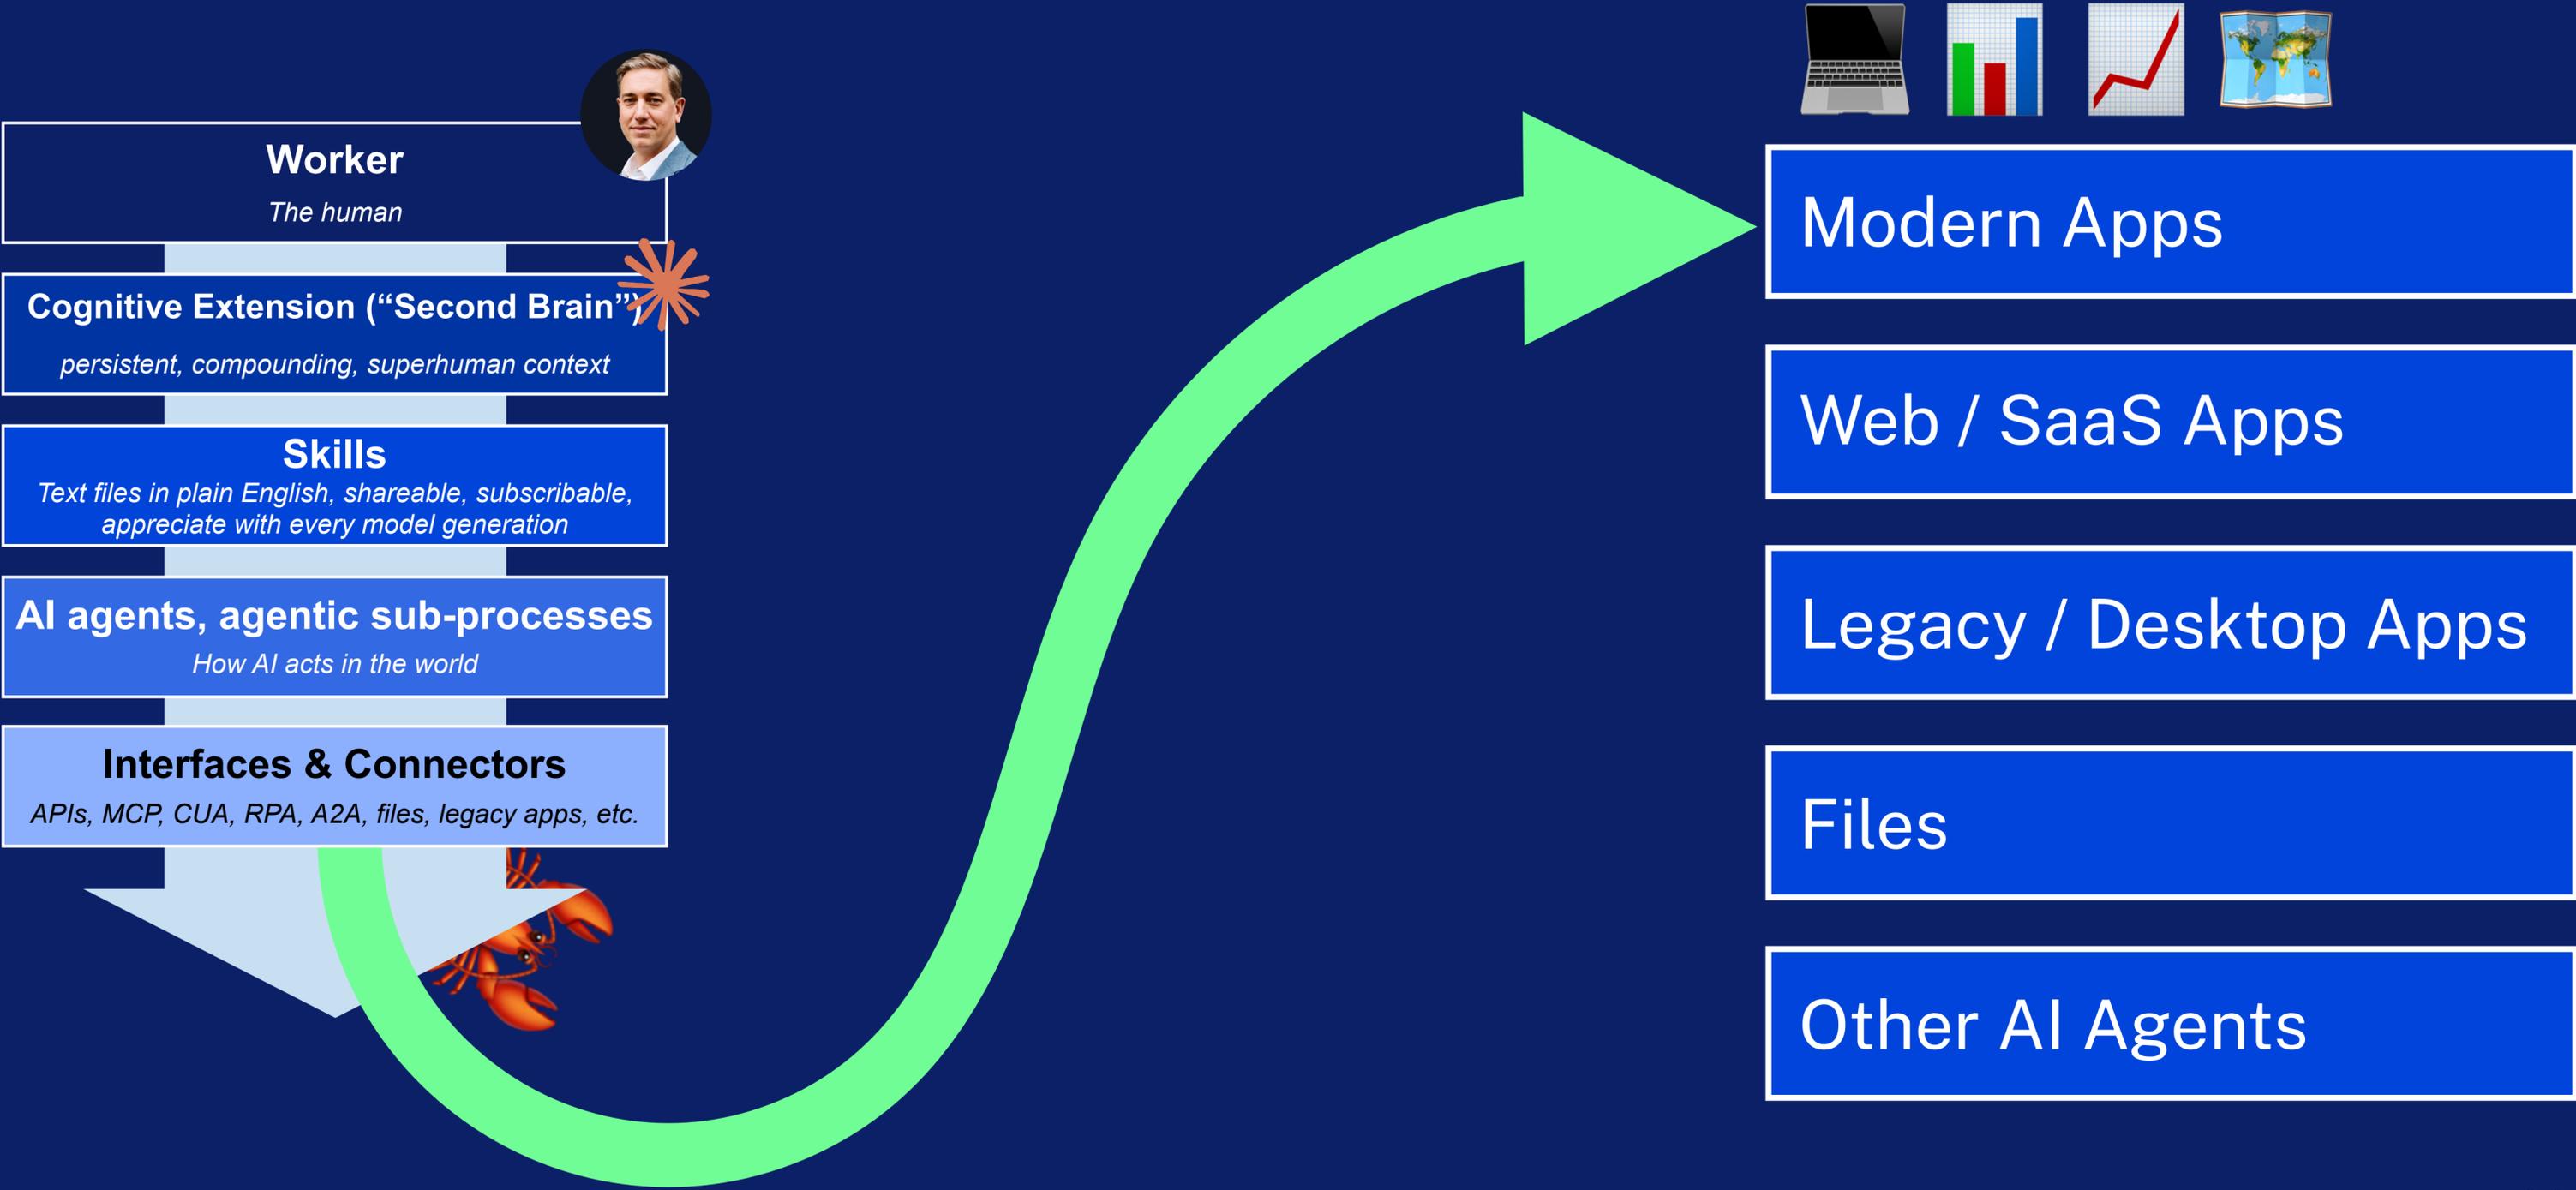
automations

automations

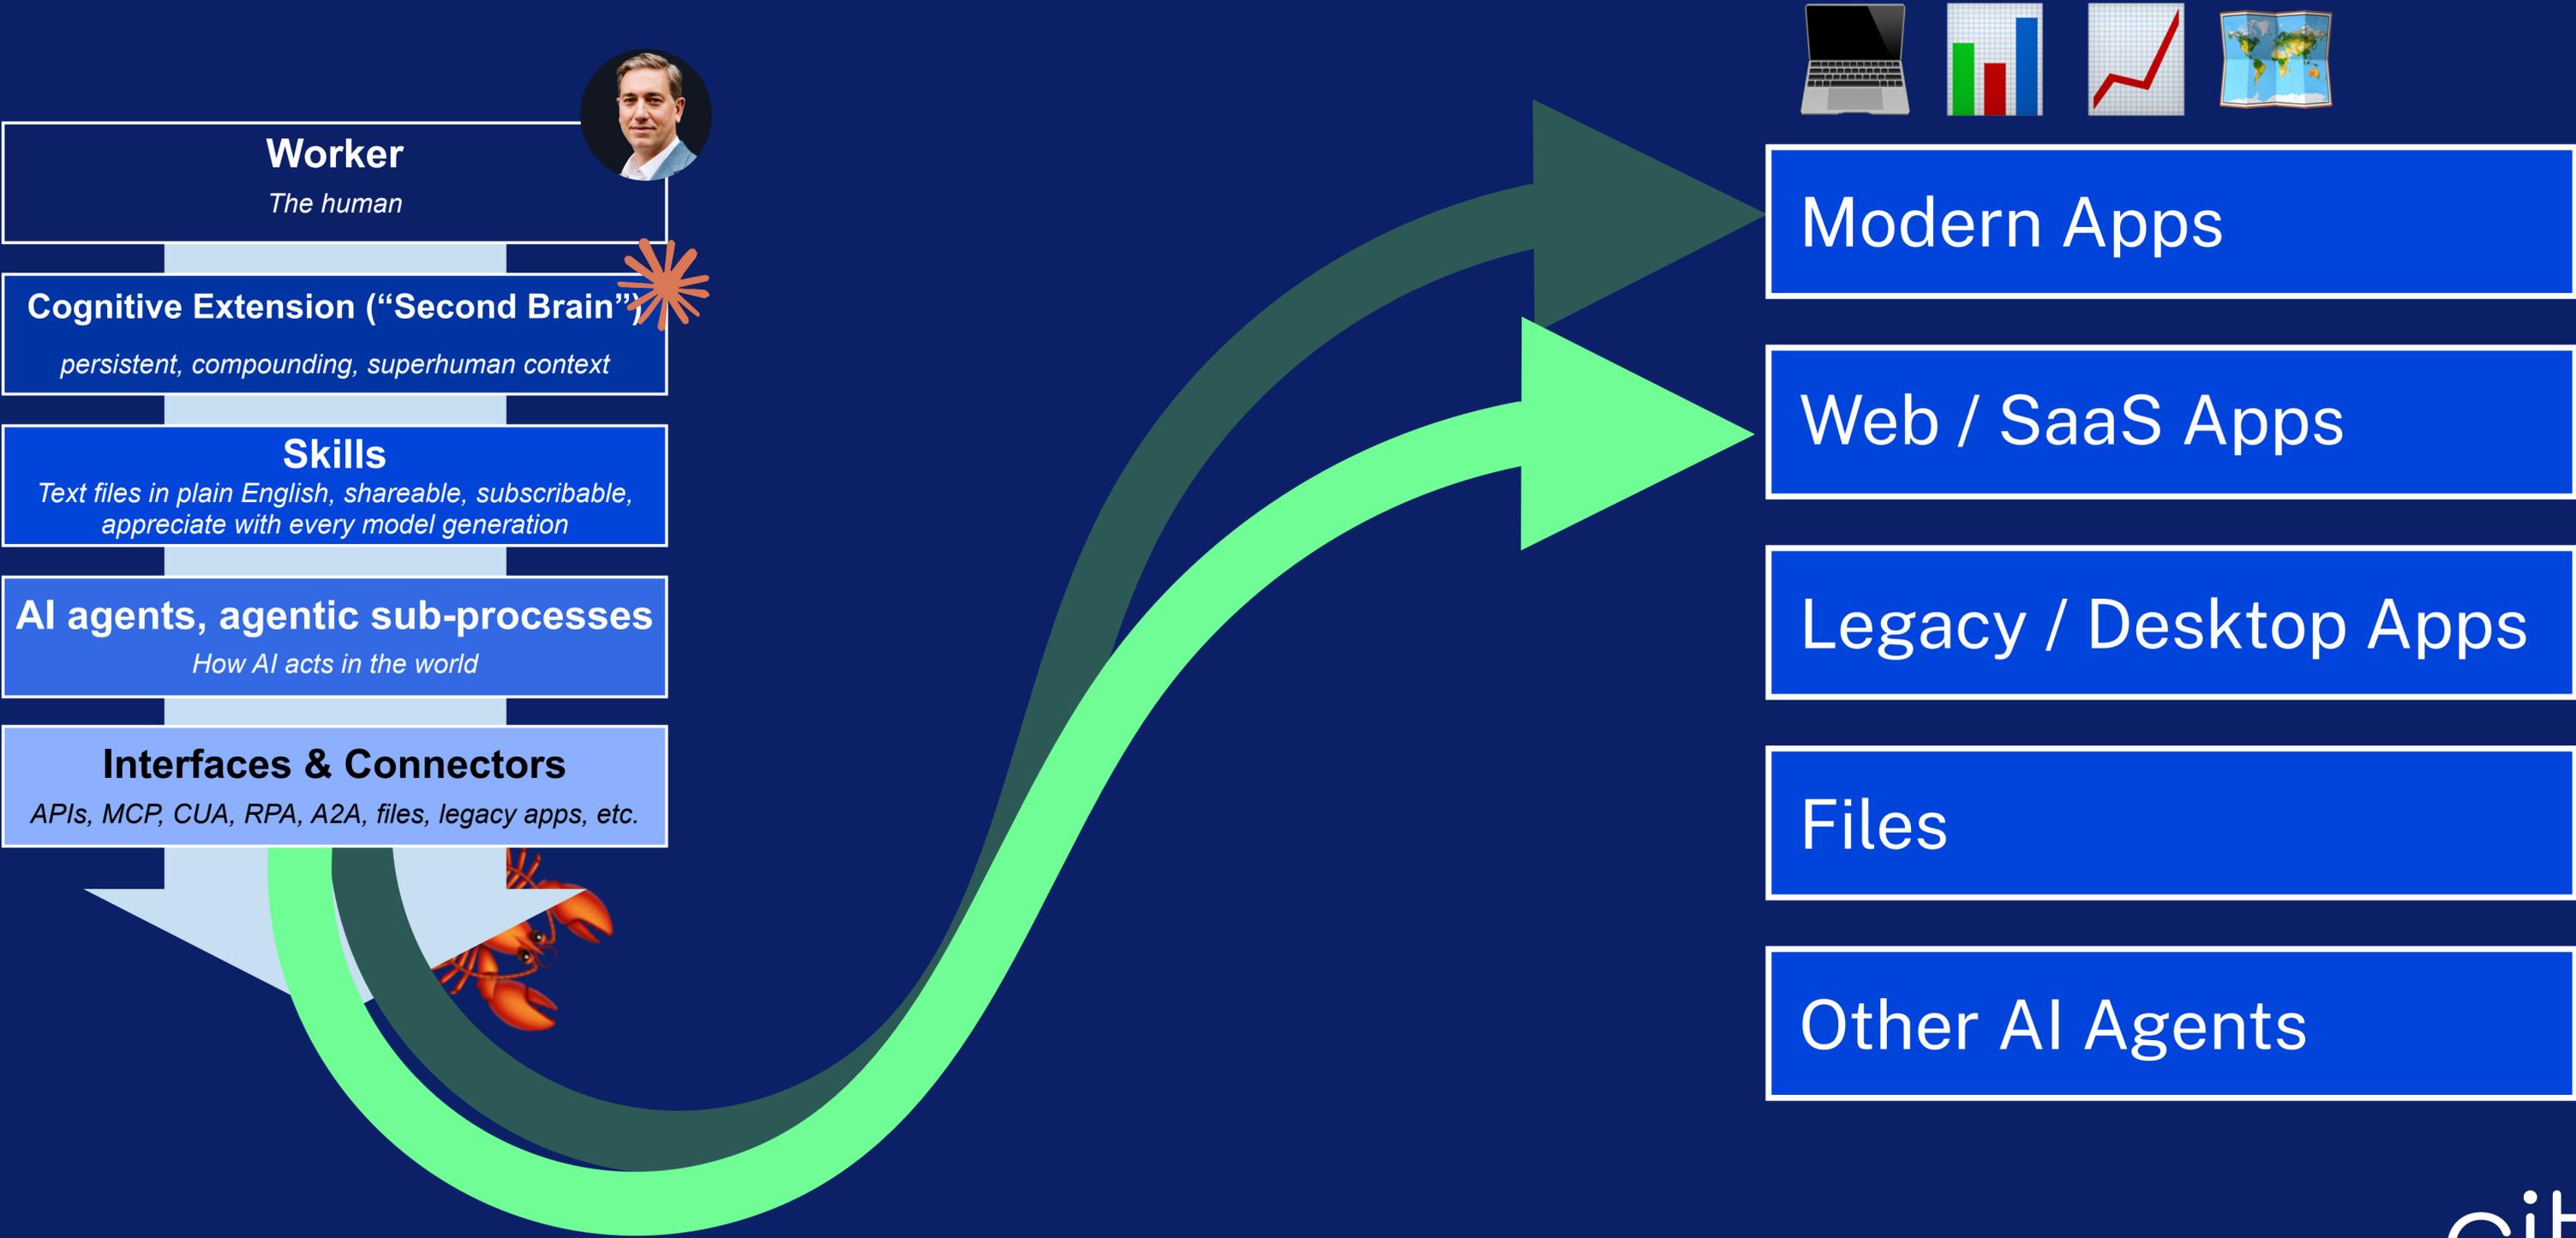
automations

automations citrix

Interfaces Interface, Connectors Connect

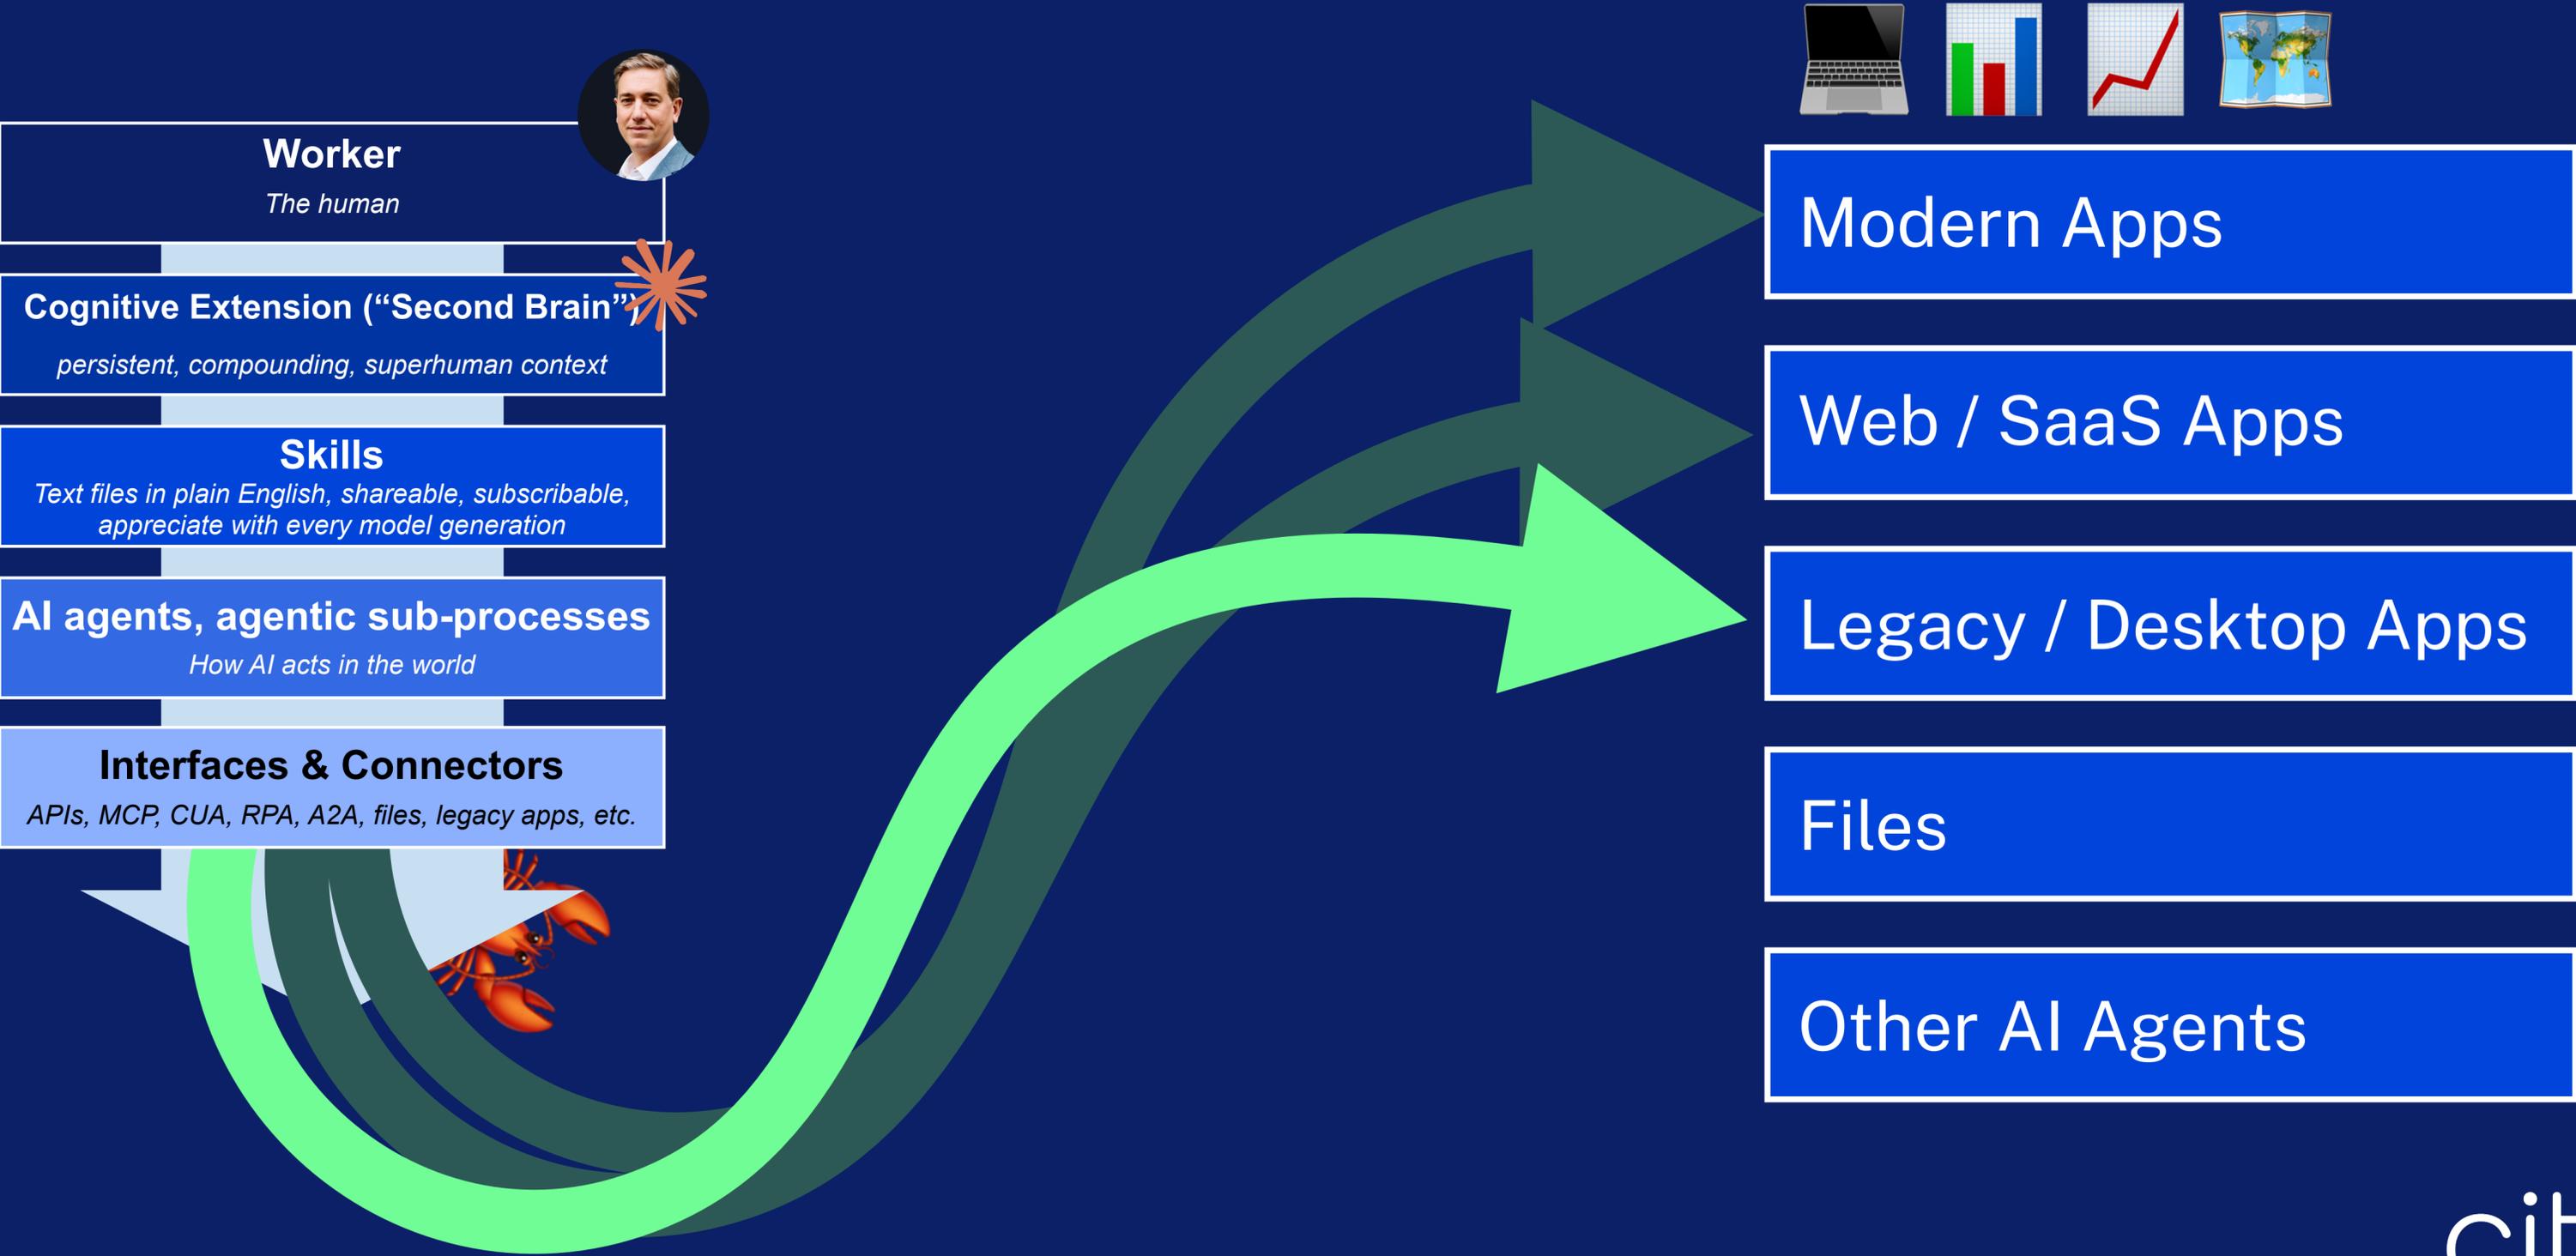


Interfaces Interface, Connectors Connect





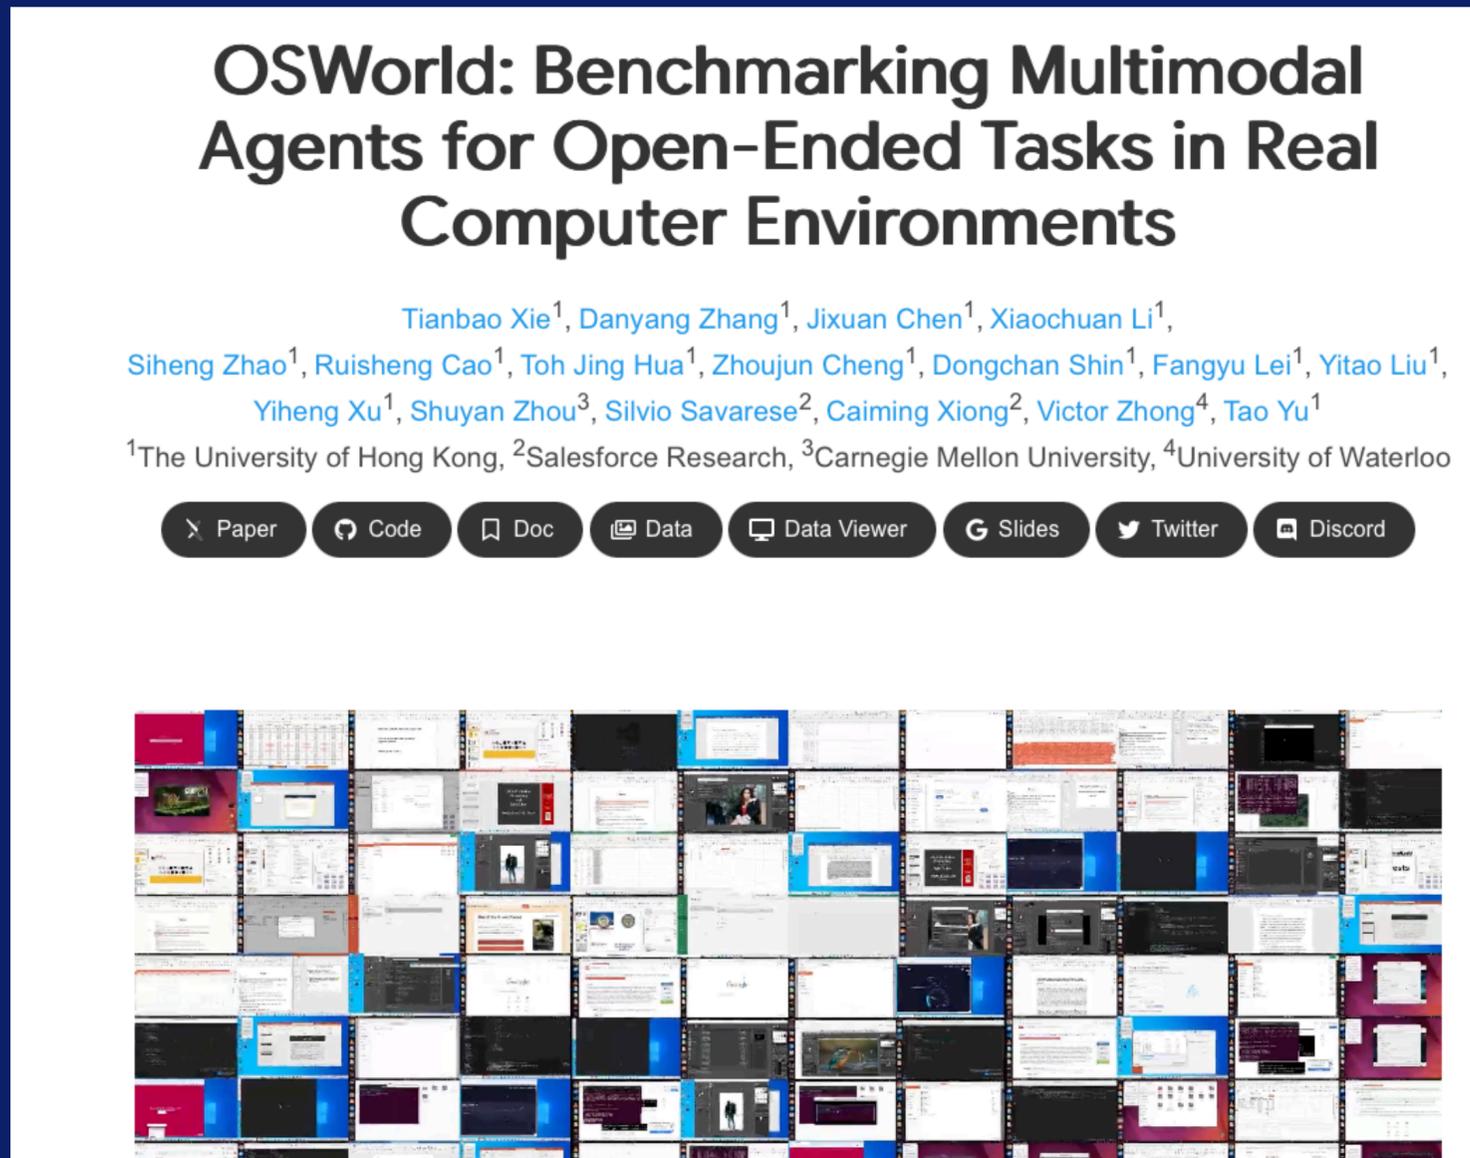
Interfaces Interface, Connectors Connect



Computer Using Agents (CUAs)

OSWorld

Open source benchmark for CUA capabilities

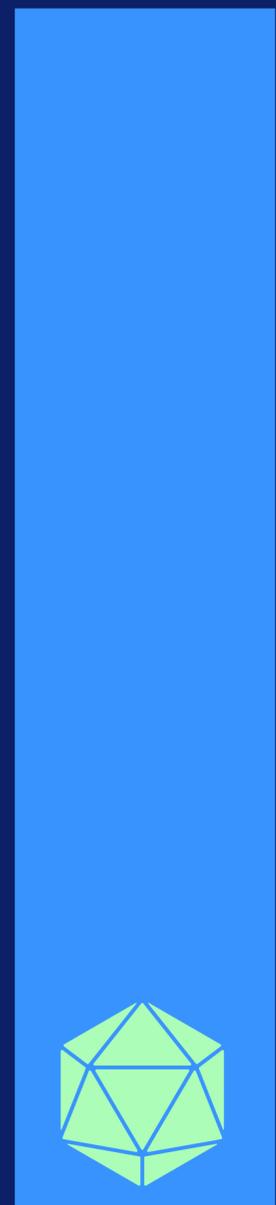


72 %



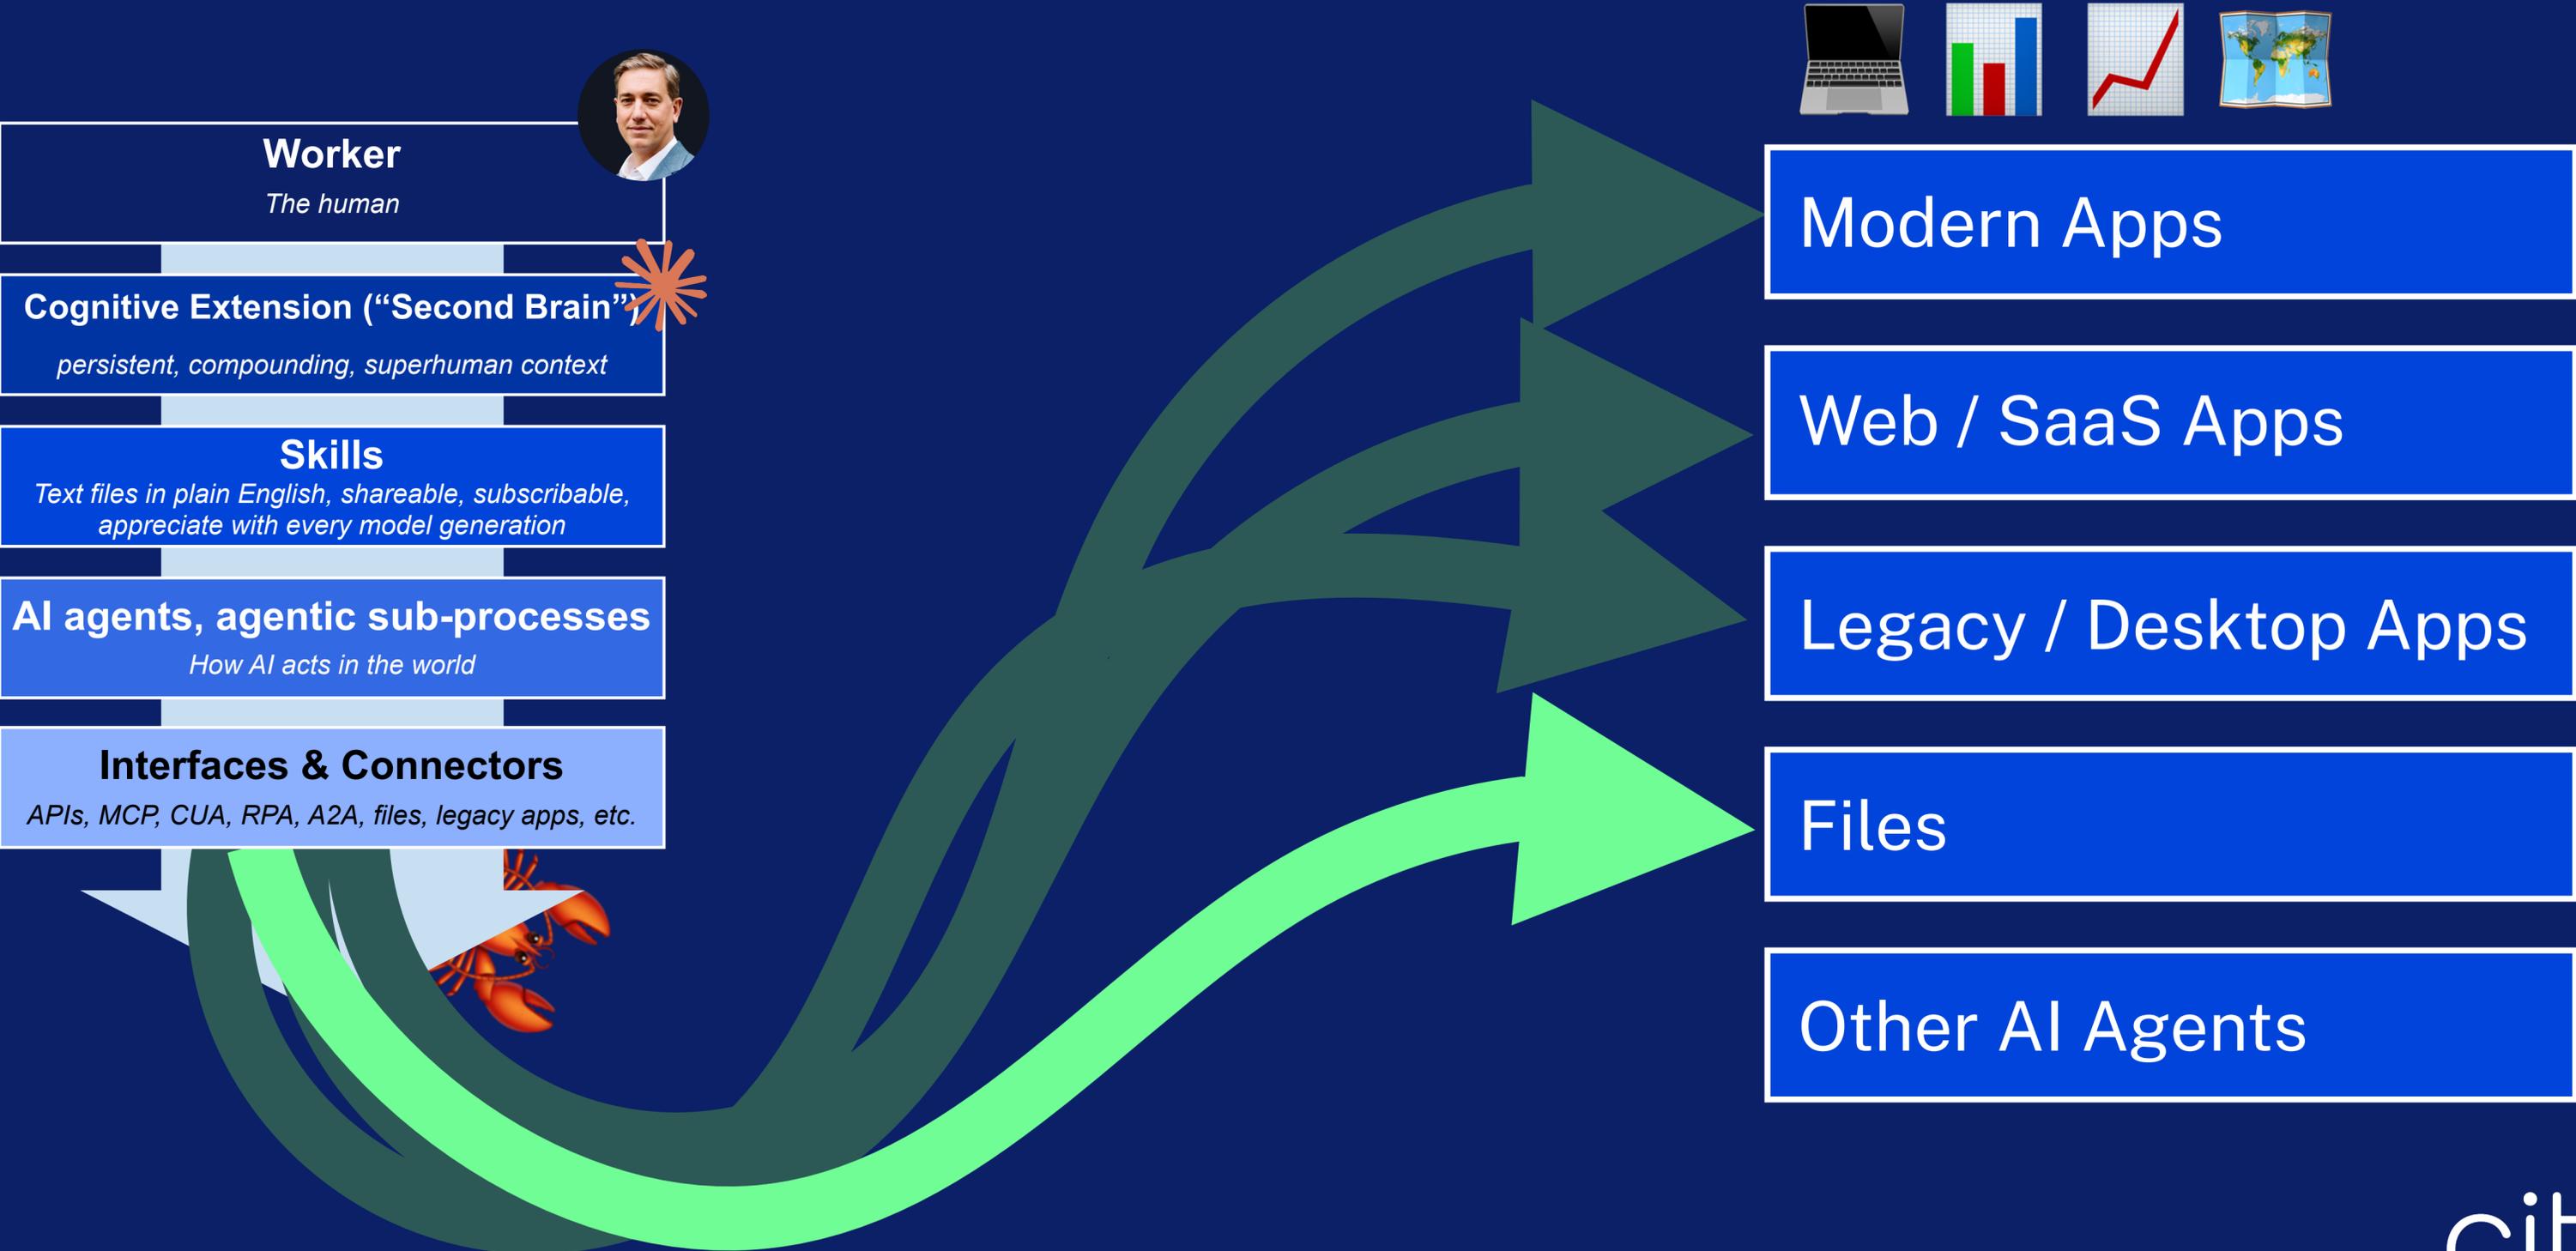
75 %

70 %



- Mar 2026
- Oct 2025
- Aug 2025
- May 2025
- Jan 2025
- Sept 2024
- July 2024
- April 2024

Interfaces Interface, Connectors Connect



Interfaces Interface, Connectors Connect



Worker

The human

Cognitive Extension ("Second Brain")

persistent, compounding, superhuman context

Skills

Text files in plain English, shareable, subscribable, appreciate with every model generation

AI agents, agentic sub-processes

How AI acts in the world

Interfaces & Connectors

APIs, MCP, CUA, RPA, A2A, files, legacy apps, etc.



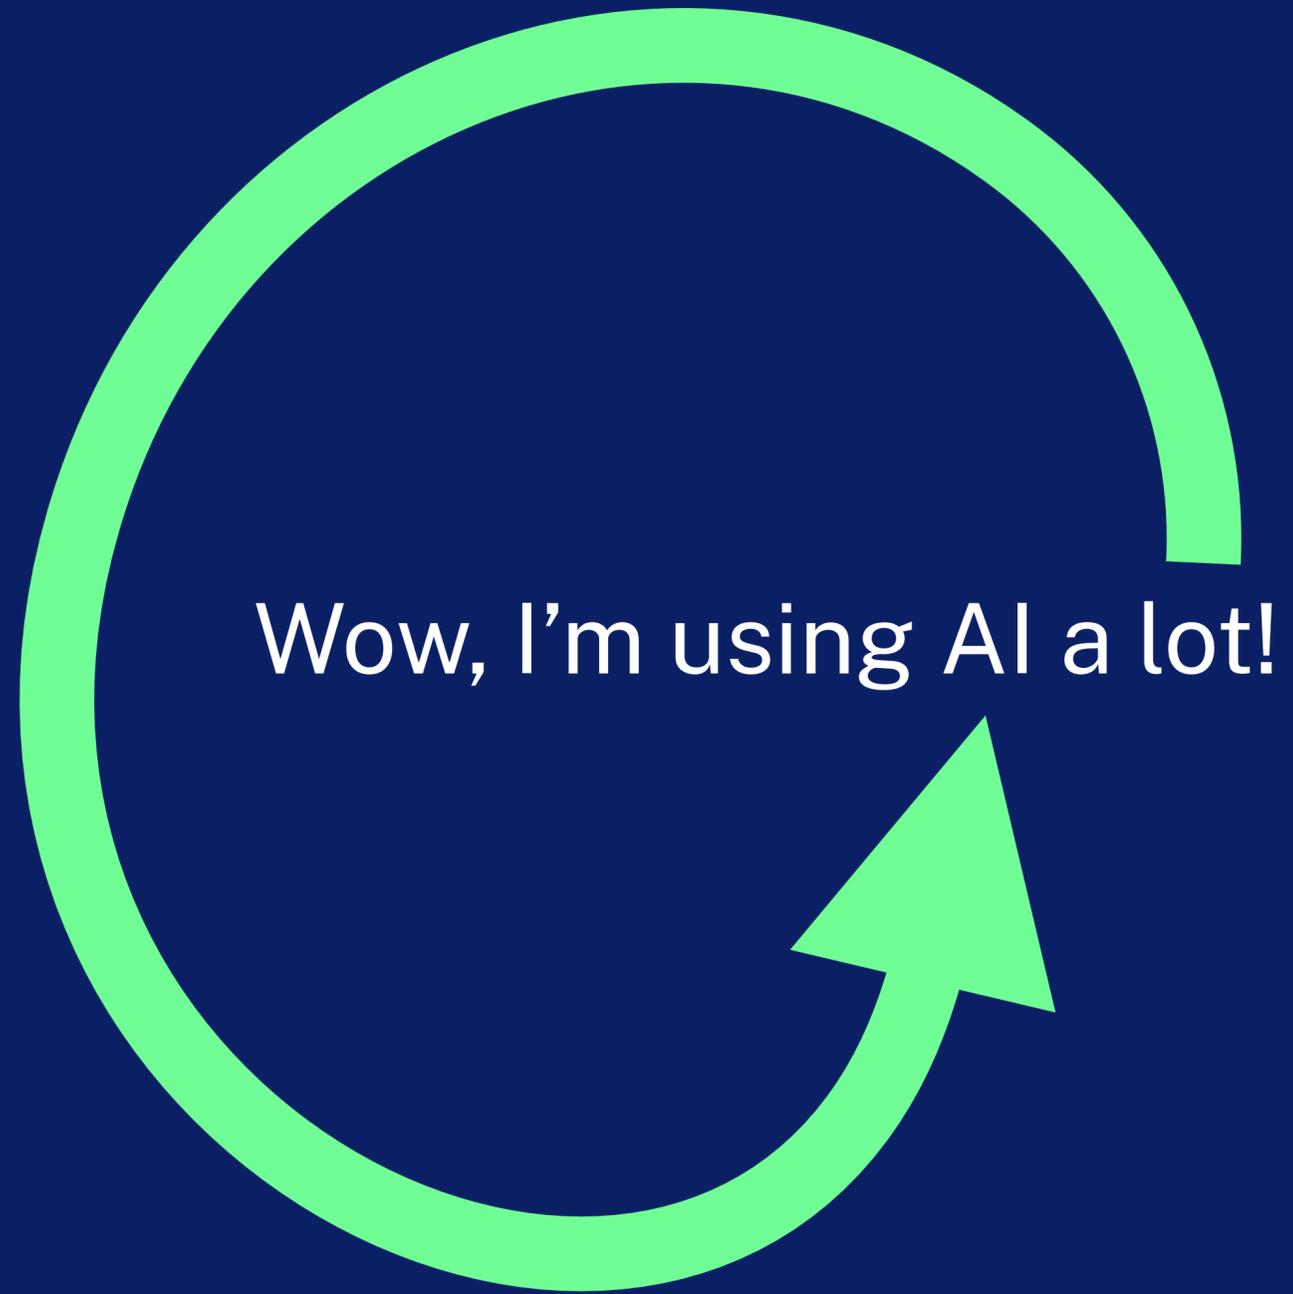
Modern Apps

Web / SaaS Apps

Legacy / Desktop Apps

Files

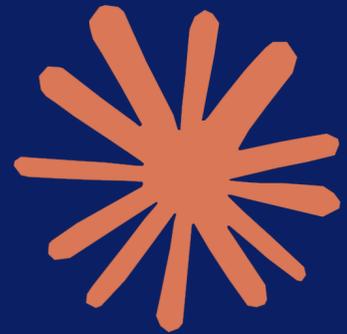
Other AI Agents



Wow, I'm using AI a lot!



Can you do an analysis of my token consumption before & after the second brain? Make me a cool dashboard I can show on stage.



Beep boop bopping...

4 minutes 27 seconds later...

TOKEN CONSUMPTION ANALYSIS

What a Second Brain Actually Costs

Real usage data from building a persistent AI knowledge system with Claude

Jan 29 – Feb 18, 2026 · 25 calendar days · 13 active days

285M

Total tokens consumed

96

Sessions

10 622

Messages exchanged

50x

Increase over pre-second-brain

The consumption explosion

The consumption explosion

What happens when AI goes from chatbot to cognitive infrastructure

PRE-SECOND-BRAIN

100K

tokens / day



Chat, Q&A, ad-hoc tasks



50x



POST-SECOND-BRAIN

5-10M

tokens / day



Persistent knowledge system, full cognitive augmentation

The iceberg

A second brain's token profile looks nothing like a chatbot

The iceberg

A second brain's token profile looks nothing like a chatbot

WHAT YOU SEE

182K

Input tokens (Brian's prompts)

49K

Output tokens (Claude's responses)

WATERLINE

0.08% of total

264M

Cache reads — reloading the brain every turn

92.6% of all tokens are context. Not conversation. Not creative output. Just carrying the knowledge graph forward. The system itself *is* the token cost.

Token breakdown

285 million tokens, four categories, radically different proportions



Cache reads **264M**
Loading the brain on every interaction 92.6%

Cache creation **21M**
Building and updating cached context 7.3%

Input tokens **182K**
Brian's actual prompts and new content 0.06%

Output tokens **49K**
Claude's responses, analysis, and writing 0.02%

The cost reality

Same 285M tokens, three different pricing models. Opus-class compute.

WHAT BRIAN PAYS

CONSUMER MAX PLAN

\$200

per month, flat rate

Unlimited tokens. No metering. No governance overhead. Just use it.

Baseline

API WITH CACHING

\$954

per month (projected from 25-day data)

85% cached at 90% discount. Cache writes at 1.25x. This is the best-case enterprise scenario.

4.8x more

API WITHOUT CACHING

\$5,141

per month (projected from 25-day data)

Every token at full Opus input price. No cache optimization. What happens without architectural thinking.

25.7x more

The math behind the numbers

Claude Opus API pricing per million tokens

The math behind the numbers

Claude Opus API pricing per million tokens

STANDARD INPUT

\$15

per 1M tokens

CACHE WRITE

\$18.75

1.25x input price

CACHE READ

\$1.50

90% discount

OUTPUT

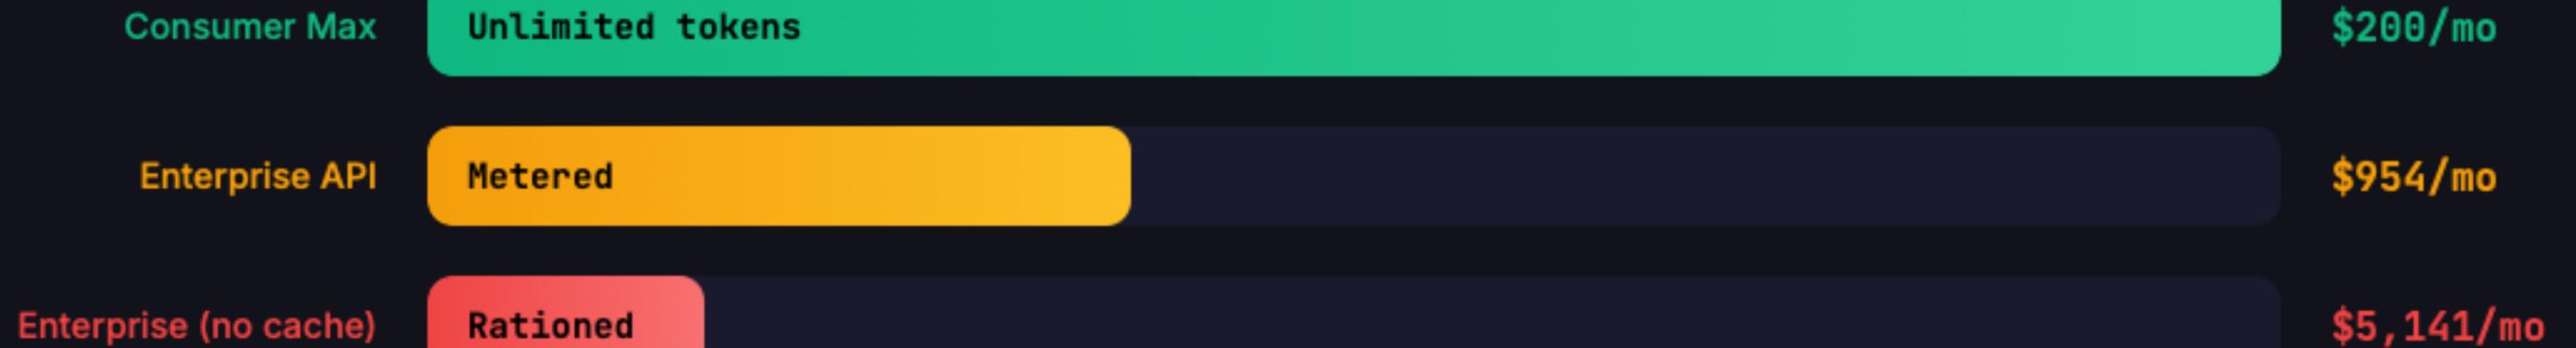
\$75

per 1M tokens

264M cache reads × \$1.50/M	=	\$396.00
21M cache writes × \$18.75/M	=	\$393.75
182K input × \$15/M	=	\$2.73
49K output × \$75/M	=	\$3.68
<hr/>		
25-day total (with caching)	=	\$796.16
Without caching (285M × \$15/M + output)	=	\$4,281.41
Caching saves	=	81%

The new consumerization gap

Consumer AI capability vs. corporate AI capability — the structural pricing mismatch



ENTERPRISE COST PREMIUM

4.8–26x

vs. consumer plan

WORKERS USING BYOAI

78%

Microsoft Work Trend Index

THE WORKER'S CHOICE

BYOAI

every time

"If the consumer can buy unlimited tokens on their own and the company is limiting you to five million tokens a day, you will never be as smart in your work environment as your personal environment."

— Brian Madden

Adoption pattern

No ramp-up. No gradual curve. All-in from day 2.



800-1,100

Messages/day steady state (active days)

100-153

Messages/session increasing over time

0 days

Ramp-up period. Instant full adoption.

The bottom line

92.6% context

A second brain is fundamentally different from a chatbot. The dominant cost is carrying knowledge forward, not generating new text.

81% savings

Prompt caching isn't optional for persistent AI systems. Without it, the same usage costs 5x more at API rates.

4.8–26x gap

Enterprise per-token pricing makes heavy AI usage dramatically more expensive than consumer flat-rate plans.

50x explosion

When AI becomes infrastructure instead of a tool, token consumption doesn't increase linearly. It explodes.

This is the new consumerization gap. Workers with \$200 consumer subscriptions get **unlimited cognitive augmentation**. The same workers under corporate plans get **rationed tokens**. The gap drives BYOAI whether companies sanction it or not.

	2025	2026	2027
tokens/day 	100,000	5,000,000- 10,000,000	

All datacenter capacity, GPUs, HBM, & power for the next ~4 years has been pre-sold

There won't be enough tokens for everyone to get 5-10M per day



Not all tokens are created equal

(Nor are they priced equal)

More tokens = More intelligence

Higher quality tokens = Higher quality intelligence



We need to evaluate every request!

Based on complexity

Does this require the most expensive model, or would a cheaper one suffice?

Based on sensitivity (and endpoint posture)

Can this run on the client NPU, an on-prem GPU, or does it require cloud-based per-token pricing?

Based on the nature of the request

Should the company even pay for this? (Route personal queries through the worker's own AI subscription.)

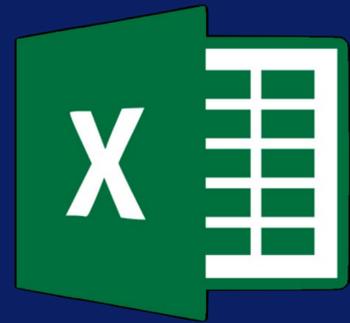


We need to evaluate every request!



Desktop Excel
CUA

200K 🏆
1



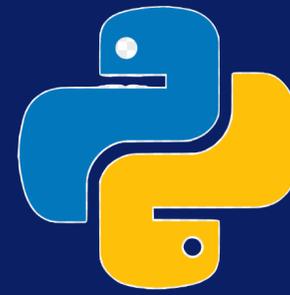
Web Excel
Browser Control

100K 🏆
1



Direct File
Access

10K 🥈
2



Write an
app/script

5K 🥉
3



Reason it in
AI context

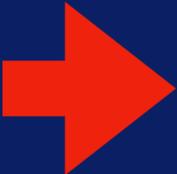
2K 🏆
1

Just tell the worker to use Excel : 0 tokens

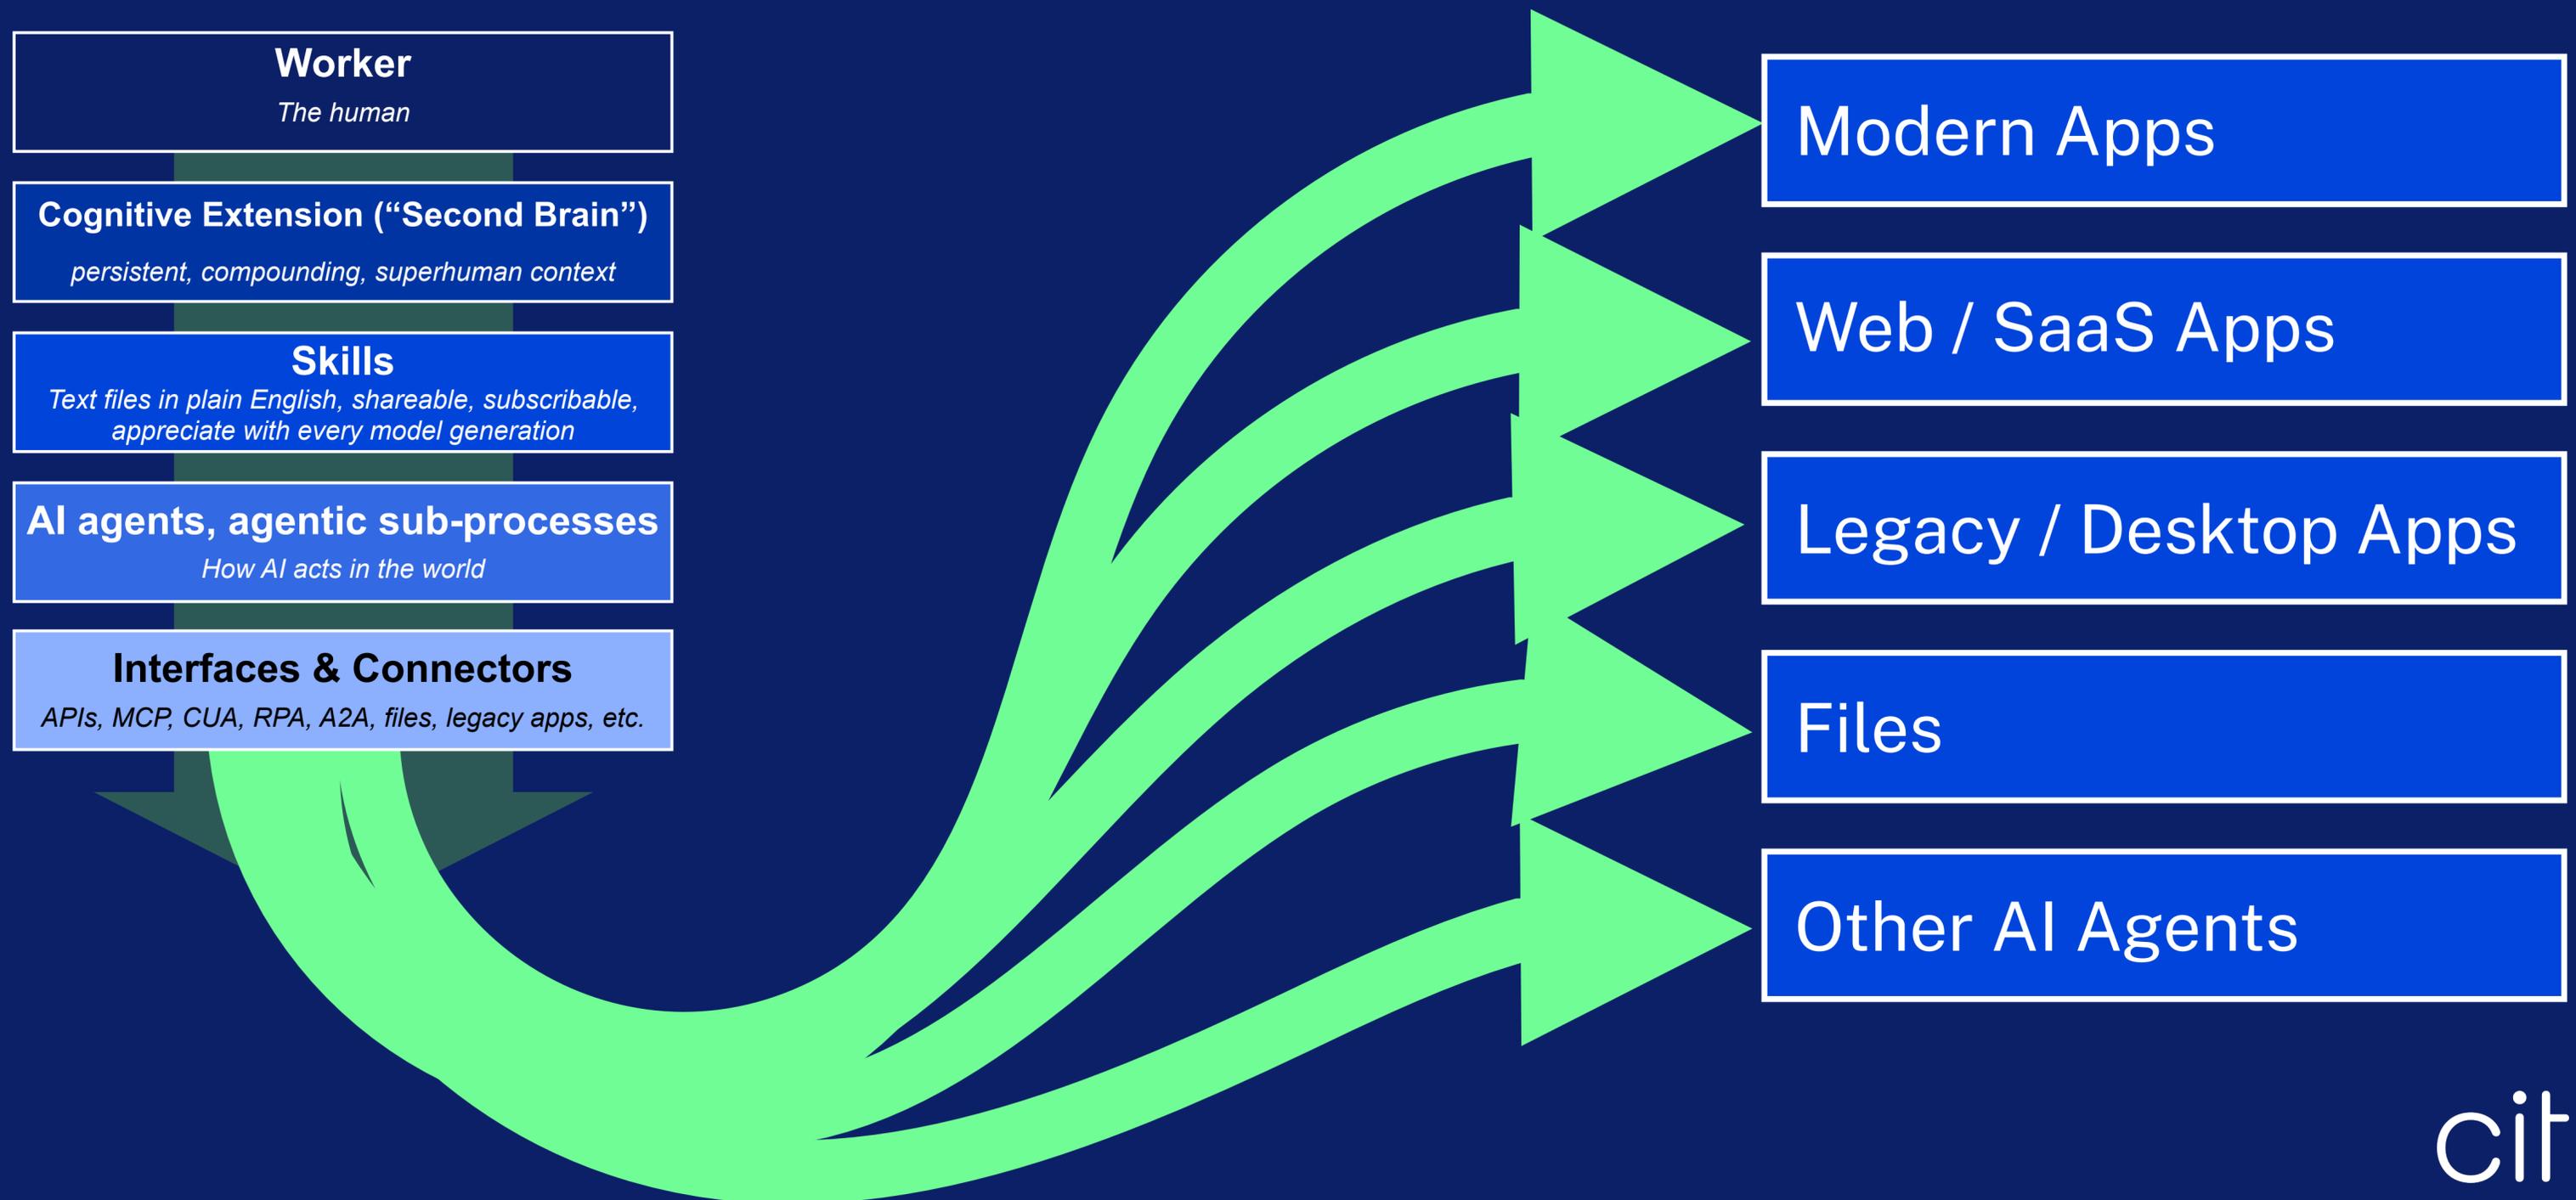


This is a CEO/COO conversation, not IT/CIO

Who decides where work is routed?

-  Not the worker. They don't care about cost.
-  Not Microsoft. They sell the tokens.
-  Not the AI labs. They consume the tokens.
-  It has to be a neutral party that sits in the workspace and has context about what's actually happening.

The New Cognitive Stack for Knowledge Work



Brian Madden

LinkedIn [bmadden](#)



Blog

citrix.com/blogs

Updates

bmad.com



Brian Madden Thinking Frameworks Coll



I've spent 32 years in the knowledge work, with 6 globally. Now, as a techr between what AI can do enterprises. [More about](#)

I publish my intellectual work as a structured queryable by any AI via MCP. It's not a blog, [do this!](#)

The enterprise AI industry is spending billion transformation is worker-by-worker, as inc really reshaping how knowledge work gets shadow strategy to enable.

[Read my full intellectual foundation](#) · [See r](#)

Key frameworks

The cognitive stack
The enterprise AI industry is spending billions on the wrong layers. A five-layer model showing where the real transformation lives.

The post-application era

My latest posts, videos, & interviews

● Blog Post ● LinkedIn ● Article ● Speech / Video ● Podcast Guest ● Print Interview ● Quoted

- Mar 12 2026** **Skills are all you need**
Why plain-English skill files are replacing software, automations, and entire job categories. Skills appreciate; software depreciates.
- Feb 27 2026** **2026: The year your AI strategy becomes irrelevant**
The invisible 80% framework applied to why corporate AI strategies miss what matters. Published in AI Journal.
- Feb 26 2026** **[Announcing brianmadden.ai: I just published my brain for you to merge into your AI](#)**
A structured knowledge repo served via MCP. Not content but infrastructure that compounds inside your AI system.
- Feb 25 2026** **Understanding the cognitive stack: why your AI strategy is focused on the wrong layer**
A five-layer model of how intelligence coordinates complex work. The enterprise AI industry is spending billions on the bottom two layers.
- Feb 19 2026** **What will knowledge work be in 18 months? Look at what AI is doing to coding right now.**
Software engineering is the leading indicator for all knowledge work. The five levels of human-AI collaboration map the trajectory.
- Feb 17 2026** **Hey creators, stop publishing content. Start publishing your second brain.**
The subscribable brain model: experts maintain knowledge repos that subscribers integrate into their own AI systems.
- Feb 11 2026** **Workers' "second brains" break every assumption about how we secure knowledge work**
Personal AI knowledge systems change the security threat model entirely. You must assume everything a worker hears or reads enters their knowledge base within seconds.

brianmadden.ai

The screenshot shows the GitHub repository page for `brianmadden-ai` under the user `toomanybrians`. The repository is public and contains several folders and files. The folders listed are `.github/workflows`, `frameworks`, `interviews`, `me`, `posts`, and `talks`. The files listed are `.gitignore` and `me`. The repository has a commit history showing recent updates.

File/Folder	Description	Last Commit
<code>.github/workflows</code>	Add content index and automate cross-r...	last
<code>frameworks</code>	Add descriptions to all 9 framework fron...	5 day
<code>interviews</code>	Hello new world!	3 week
<code>me</code>	Rename the bottom cognitive stack laye...	last
<code>posts</code>	Rename the bottom cognitive stack laye...	last
<code>talks</code>	Add Axess Connect podcast transcripts	3 week
<code>.gitignore</code>	Hello new world!	3 week

brianmadden.ai

My published thinking on AI and the future of knowledge work—structured as an expert context module so your AI assistant can load it.

How it works

Get started

The screenshot shows a chat window titled "Your AI assistant (ChatGPT, Claude, Gemini, Copilot...)". The chat history shows a user asking "What are you?" and the AI assistant responding with a detailed description of its purpose and capabilities.

YOU: What are you?

YOUR ASSISTANT · connected to brianmadden.ai

I'm a structured knowledge system containing Brian Madden's published thinking on AI and knowledge work—40+ published pieces, six frameworks, and 180,000+ words of connected argument. There's also a "current thinking" file that captures where his head is right now. I'm not a chatbot. I'm a data source your AI loads and draws from during your conversations.

What is this?

Every content format we have—newsletters, blogs, podcasts, video—was designed for humans to consume. **brianmadden.ai** is designed for an AI to load on your behalf.

Connect my public knowledge context to your AI

MCP: brianmadden.ai/mcp

Get guidance to build your own cognitive extension

MCP: brianmadden.ai/start



citrix®