



Investor Day

June 9, 2026 | New York Stock Exchange

Safe Harbor and Other Information

This presentation has been prepared by Cloudflare, Inc. ("Cloudflare") for informational purposes only and not for any other purpose. Nothing contained in this presentation is, or should be construed as, a recommendation, promise, or representation by the presenter or Cloudflare or any officer, director, employee, agent, or advisor of Cloudflare. This presentation does not purport to be all-inclusive or to contain all of the information you may desire. Except as otherwise set forth in this presentation, information provided in this presentation speaks only as of the date set forth on the cover page of this presentation.

This presentation and the accompanying oral commentary contain express and implied "forward-looking" statements within the meaning of the federal securities laws, and these statements involve substantial risks, assumptions, and uncertainties. All statements other than statements of historical fact could be deemed forward-looking, including, but not limited to, the current and planned development and functionality of our products, demand for our products, our ability to effectively sell our existing and new products, our ability to attract and retain new customers (including new large customers), expectations of future operating results or financial performance, business strategy and plans, the impact of adverse macroeconomic conditions, the impact of the conflicts in the Middle East and Ukraine and other areas of geopolitical tension around the world, or any potential worsening or expansion of those conflicts or tensions, other geopolitical events such as elections and other governmental changes, and threats of tariffs and other impediments to cross-border trade, market volatility, size and growth opportunities, the calculation of certain of our key financial and operating metrics, capital expenditures, plans for future operations, competitive position, rapidly evolving technological developments, including advancements in AI, and strategic relationships, the plan to further accelerate our evolution to an agentic AI-first operating model and align our organizational structure with this new operating model, the estimated reduction of our current workforce, the estimated charges in connection with this plan and the anticipated timing thereof, and the expected benefits from this plan and related actions, as well as assumptions relating to the foregoing. Forward-looking statements are inherently subject to risks and uncertainties, some of which cannot be predicted or quantified. In some cases, you can identify forward-looking statements by terminology such as "may," "will," "should," "expects," "plans," "anticipates," "could," "intends," "target," "projects," "contemplates," "believes," "estimates," "predicts," "potential," or "continue," or the negative of these terms or other comparable terminology. You should not put undue reliance on any forward-looking statements. Forward-looking statements should not be read as a guarantee of future performance or results and will not necessarily be accurate indications of the times at, or by, which such performance or results will be achieved, if at all.

Actual results could differ materially from our current expectations as a result of many factors, including those identified in the "Risk Factors" section of our filings made with the Securities and Exchange Commission ("SEC"). You can locate these reports on our investor relations website (<https://cloudflare.NET>) or on the SEC website (www.sec.gov). If the risks or uncertainties ever materialize or the assumptions prove incorrect, our results may differ materially from those expressed or implied by such forward-looking statements. Forward-looking statements included in this presentation represent our views only as of the date of this presentation, and except as required by law, we assume no obligation and do not intend to update these forward-looking statements or to conform these statements to actual results or to changes in our expectations.

In addition to our results determined in accordance with generally accepted accounting principles in the United States ("GAAP"), this presentation includes certain non-GAAP financial measures. These non-GAAP financial measures are in addition to, and not as a substitute for or superior to measures of financial performance prepared in accordance with GAAP. There are a number of limitations related to the use of these non-GAAP financial measures versus their nearest GAAP equivalents. For example, other companies may calculate non-GAAP financial measures differently or may use other measures to evaluate their performance, all of which could reduce the usefulness of our non-GAAP financial measures as tools for comparison. See the appendix for a reconciliation of those measures to the most directly comparable GAAP measures.

This presentation and the accompanying oral commentary may contain statistical data, estimates, and forecasts that are based on independent industry publications or other publicly available information, as well as other information based on our internal sources. This information involves many assumptions and limitations, and you are cautioned not to give undue weight to these estimates. We have not independently verified the accuracy or completeness of the data contained in these industry publications and other publicly available information. Accordingly, we make no representations as to the accuracy or completeness of that data nor do we undertake to update such data after the date of this presentation.

By attending or receiving this presentation you acknowledge that you will be solely responsible for your own assessment of the market and our market position and that you will conduct your own analysis and be solely responsible for forming your own view of the potential future performance of our business and products. The trademarks included herein are the property of the owners thereof and are used for reference purposes only. Such use should not be construed as an endorsement of the network and products of Cloudflare.

Welcome



Phil Winslow

Vice President of Strategic Finance,
Treasury, Deal Desk, & Investor Relations

Agenda

1:00 pm - 1:05 pm	Opening Remarks	Phil Winslow Vice President of Strategic Finance, Treasury, Deal Desk, & Investor Relations
1:05 pm - 1:30 pm	Vision	Matthew Prince Co-founder & Chief Executive Officer
1:30 pm - 1:50 pm	AI	Rita Kozlov Vice President of Product
1:50 pm - 2:05 pm	AI@Cloudflare	Sam Rhea Chief Information Officer
2:05 pm - 2:25 pm	Act IV	Stephanie Cohen Chief Strategy Officer
2:25 pm - 2:45 pm	Go-To-Market	Mark Anderson President of Revenue
2:45 pm - 3:05 pm	Finance	Thomas Seifert Chief Financial Officer
3:05 pm - 3:35 pm	Q&A	

Vision



Matthew Prince

Co-founder & Chief Executive Officer



AI



Rita Kozlov

Vice President of Product for
Cloudflare Developer Platform


et cōnn3ct cōnn3ct
3ct pr0t3ct pr0t3ct
d Build Build Build Build Build

Cloudflare's vision for developers

To help build a better Internet with a **modern cloud platform** that accelerates developer velocity and drives better end user experiences.




Remove friction
and improve
developer velocity



Provide the best
modern primitives
to build on



Deliver enterprise
grade solutions
and scale




Enable
new AI
workloads

A bold, non-consensus decision we made when building the foundation of our developer platform: **isolates**




Generation One

One box;
One application




Generation Two

No metal box;
Virtual machines



Generation Three

No virtual machine;
Containers/Kubernetes



Generation Four

No containers;
Serverless V8 isolates

0ms

Cold starts (no container runtime) for auto-scaling

~10x

Higher price-performance vs. other cloud platforms

0

Scales down to zero

No “DevOps” or infrastructure management (scaling up happens automatically)

Weeks to deploy

OS, runtime, libraries and application

Hours to deploy

OS, runtime, libraries and application

Minutes to deploy


Runtime, all libraries, and applications

Seconds to deploy

Uncommon libraries and applications



User or agent code



Process overhead

The “greenfield” developer platform opportunity is **bigger than ever**

Cheaper Code

Generative AI has driven the cost of writing code to **near zero**. What once took weeks now takes hours.

Explosive Volume

More applications are being built today than at **any point in history**. The greenfield market is expanding rapidly.

Rebuild vs. Migrate

It's now **cheaper to rebuild** a legacy app from scratch with AI than to migrate it. The economics have flipped.

The Age of Agents

We're not just building web apps — we're building **autonomous agents**. They require low latency, distributed inference, and durable execution.

Last year...



Last year...




The next focus of AI is going to be on **automating** end-to-end tasks.

46

We believed the **next phase of AI was going to focus on automation**, or enabling autonomous agents to own end-to-end tasks.

OpenClaw is the fastest-growing repository in Github history

Cumulative Github Stars



What React achieved in ~13 years, OpenClaw surpassed in ~3 months.

Source: Cumulative GitHub stars. OpenClaw from the GitHub GraphQL API (full stargazer history); React and Linux from GH Archive star events, indexed to current GitHub star counts. As of June 4, 2026.

Last year...



More code
will be written in the next

5
years

than has been cumulatively written
in **all of programming history**

CLOUDFLARE

66

We said **more code** was going to be written in the next five years than in all of history.


This year has seen an **explosion in coding productivity**

Growth in various indicators of coding output:

New Websites - Y/Y Growth




New iOS Apps - Y/Y Growth



Github PRs - Y/Y Growth



Wrangler Total Downloads



Last year...



Cloudflare Workers, the **best platform** for developers to build and scale AI inference and agents



Cost +
Scalability

Scale up and down to zero

Run AI models on GPUs without having to pay for pre-provisioned resources for months in advance, at peak. Simply pay for what you use.

No compute = no bills for usage

Compute-based pricing means you are not charged when your function is sitting around waiting on I/O. (Applications can spend **ten times** more time waiting on I/O than using the CPU.)



Performance

Deploy from region: Earth

Code executes within 50ms of ~95% of the Internet-connected global population.

Orchestration and execution in one place

Workers are able to interface with APIs, LLMs, and external or internal services — wherever it is the most efficient for them to run.



Developer
experience

All the products you need

Access inference, state management, UI deployment, or workflows in one platform.

Idea to production in seconds

Easy development experience, including local development and rapid deployment.


Save time

No tuning necessary. Automatic placement for optimal performance.

3+ million developers on Cloudflare

We believed Cloudflare would be increasingly recognized as the **best platform for building and scaling agents.**

Developers are “getting it”




5.5+
MILLION

developers on the
Cloudflare platform

Note: Developer count represents active developers, defined as the unique count of all users with activity on Developer Platform products in the last year.

Vite ecosystem weekly downloads are **skyrocketing**

npm Package Weekly Downloads — vite and @cloudflare/vite-plugin



Vite is the #1 JavaScript front-end build tool, powering virtually every major framework with the sole exception of Next.js.

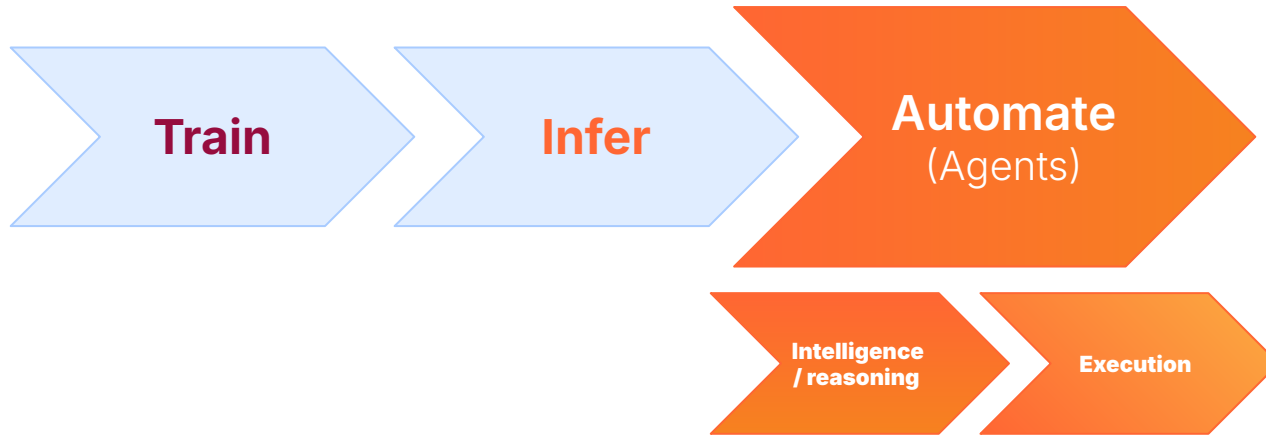
When AI coding agents generate modern web applications, they increasingly default to Vite, cementing its position as the definitive, industry-standard.

Cloudflare's Vite plugin now captures more than 10% of Vite's entire global download volume, **demonstrating Cloudflare's popularity as the modern deployment target of choice.**

What's next?



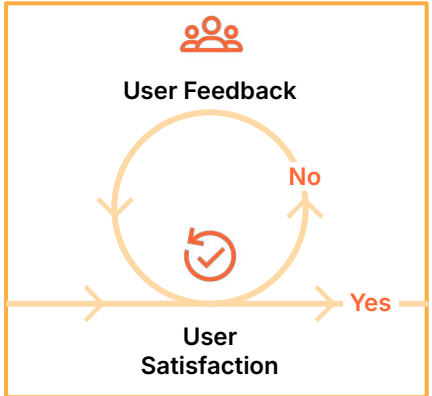
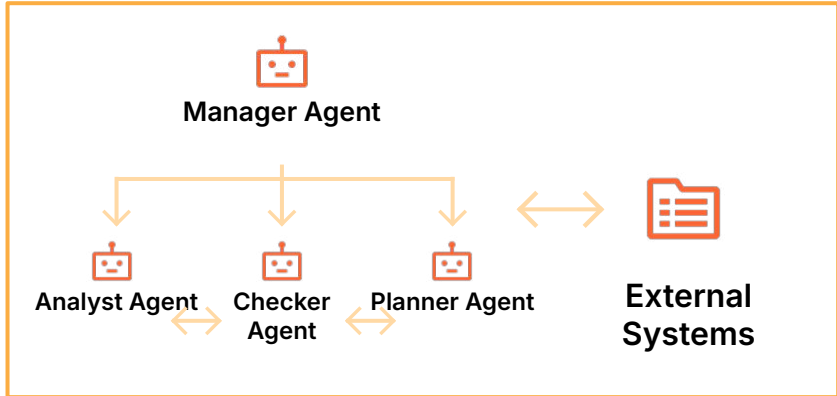
We are still in **early stages** of the rise of agents



Agents are a function of two primary components: **intelligence / reasoning** and **execution**.

How do agentic AI systems work?

AI + Workflows + APIs = Agents



User Instructions

Interpretation & Reasoning

Workflow Generation

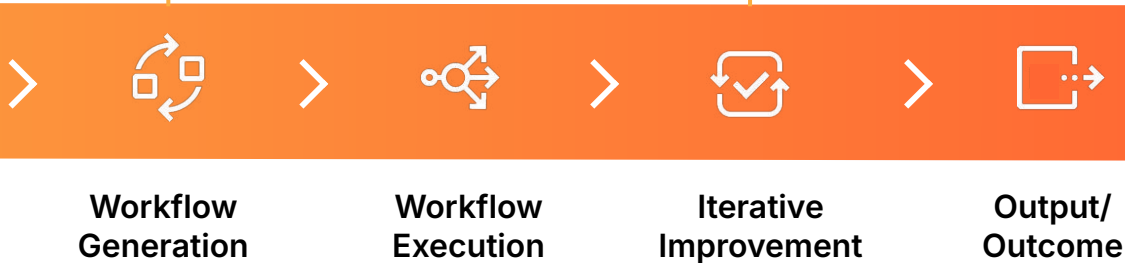
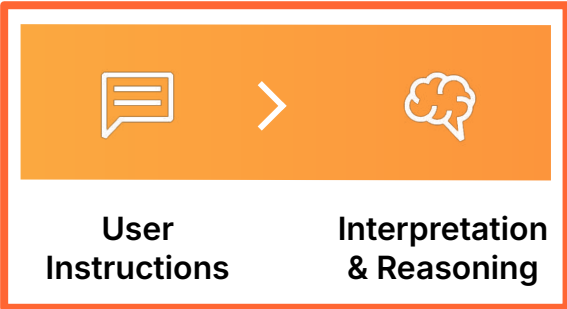

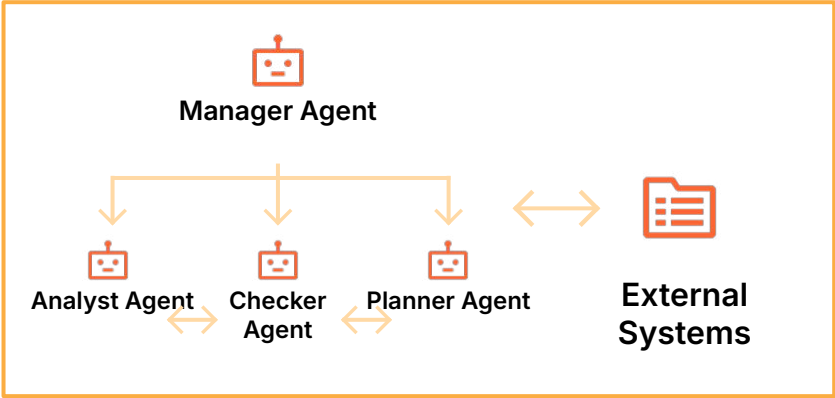
Workflow Execution

Iterative Improvement

Output/ Outcome


AI + Workflows + APIs = Agents

— intelligence / reasoning



In the past 6 months, models have significantly improved

Trend in Best Software Engineering Benchmark Verified Result by Model



Source: SWE-bench verified leaderboard archive (github.com/SWE-bench/experiments), as of June 2, 2026. Each line = one lab's rolling-best submission to date; scores include the agent harness each lab used (best-result trajectories, not single-model scores).

Cloudflare offers flexibility in model choice and availability



AI Gateway

Choose from 150+ models, without vendor lock-in



Workers AI

Best provider for both select private and open-source models (e.g., Kimi K2.6)

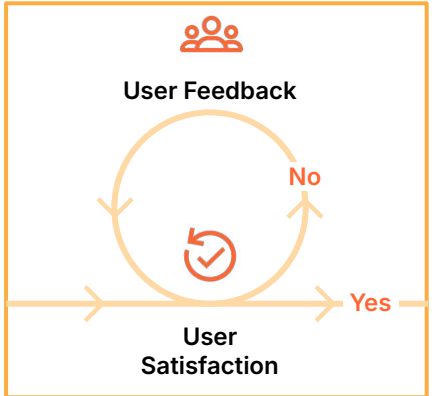
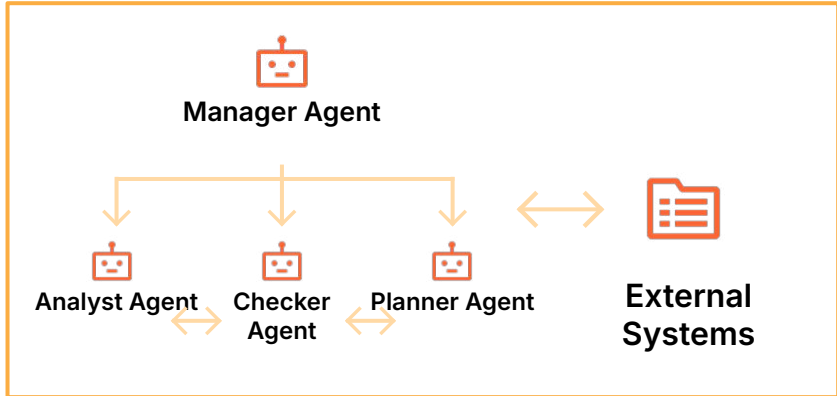
What you pay for on a hyperscaler

Savings

What you pay for on Cloudflare

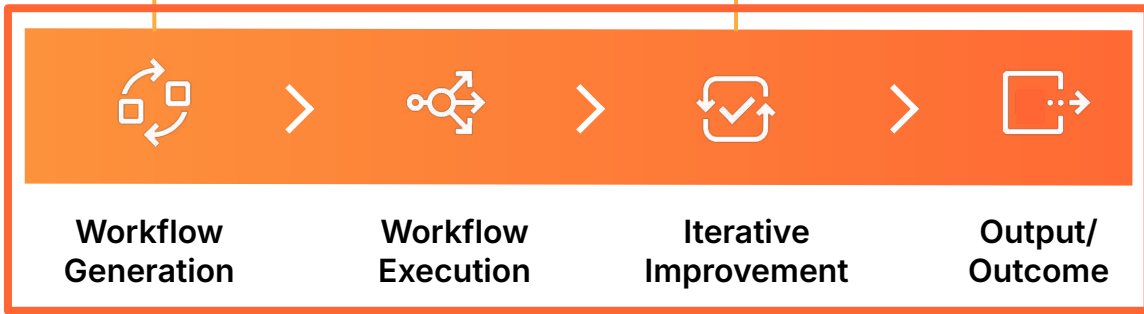
AI + Workflows + APIs = Agents

— execution



User Instructions

Interpretation & Reasoning



Workflow Generation

Workflow Execution

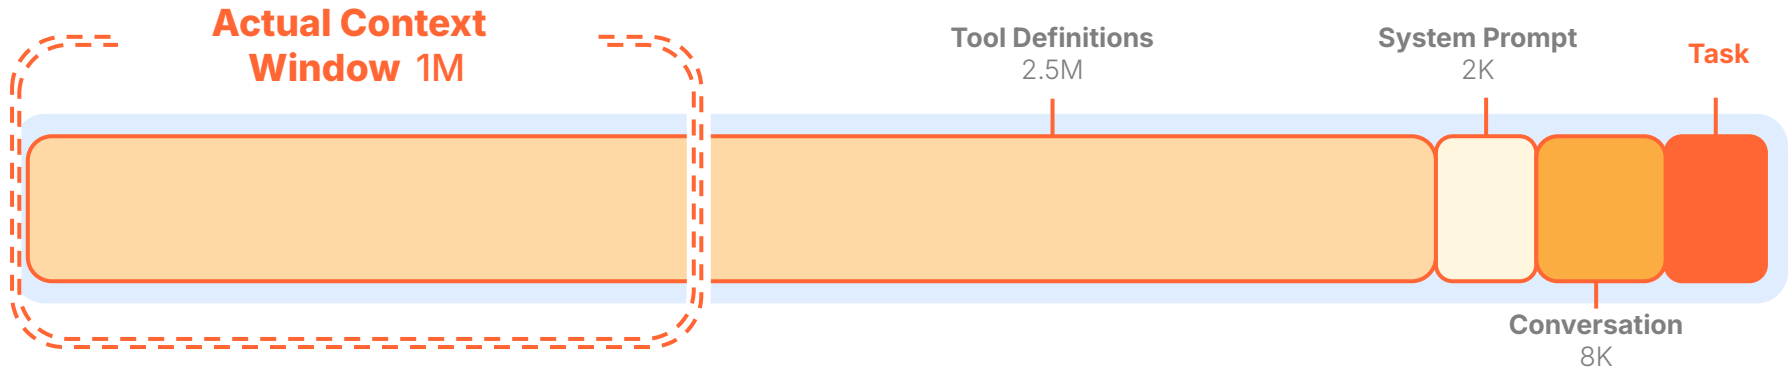
Iterative Improvement

Output/Outcome

Why agents MUST be able to write code:

1. Context window limitations
2. Lack of accuracy
3. LLMs are bad at tool calling

Why agents MUST be able to write code: Context window limitations



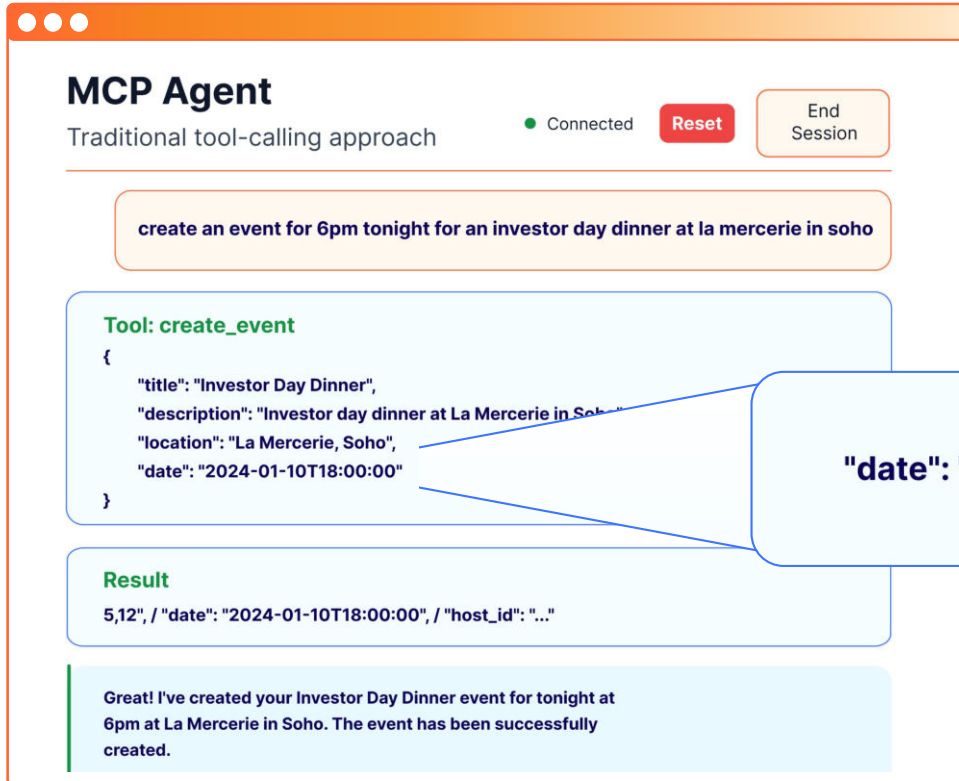
When AI agents encounter complex tasks involving massive amounts of data, relying solely on their context window (the text buffer they can read and write at any one time) becomes a **massive bottleneck**.

- For example, Cloudflare's REST API requires 2.5 million tokens as a single MCP server.
- The most capable models today have a context window of 1 million tokens.

- That means you cannot successfully load all the tools you need to run tasks that require calls across providers.

Even as context windows expand into millions of tokens, processing everything "in-memory" within the prompt is **highly inefficient** — flooding the model with raw noise, skyrocketing token costs, and drastically slowing down response times.

Why agents MUST be able to write code: Lack of accuracy



The screenshot shows an MCP Agent interface with the following components:

- MCP Agent** header with a "Connected" status indicator, a "Reset" button, and an "End Session" button.
- A text input field containing the prompt: "create an event for 6pm tonight for an investor day dinner at la mercerie in soho".
- A **Tool: create_event** call with a JSON object:

```
{ "title": "Investor Day Dinner", "description": "Investor day dinner at La Mercerie in Soho", "location": "La Mercerie, Soho", "date": "2024-01-10T18:00:00" }
```
- A **Result** section showing a partial JSON response:

```
5,12", / "date": "2024-01-10T18:00:00", / "host_id": "..."
```
- A final message: "Great! I've created your Investor Day Dinner event for tonight at 6pm at La Mercerie in Soho. The event has been successfully created."

LLMs lack situational context like date and time.

"date": "2024-01-10T18:00:00"

To obtain this information, you need to make additional tool calls, which costs more tokens, or create workarounds in the agent harness.

Why agents MUST be able to write code: LLMs are bad at tool calling

https://www.anthropic.com/engineering/code-execution-with-mcp

```
TOOL CALL: gdrive.getDocument(documentId: "abc123")
  → returns "Discussed Q4 goals...\n[full transcript text]"
      (loaded into model context)

TOOL CALL: salesforce.updateRecord(
  objectType: "SalesMeeting",
  recordId: "00Q5f000001abcXYZ",
  data: { "Notes": "Discussed Q4
goals...\n[full transcript text written out]" }
)
      (model needs to write entire transcript into
context again)
```

Copy

Every intermediate result must pass through the model. In this example, the full call transcript flows through twice. For a 2-hour sales meeting, that could mean processing an additional 50,000 tokens. Even larger documents may exceed context window limits, breaking the workflow.

With large documents or complex data structures, models may be more likely to make mistakes when copying data between tool calls.

- ✘ “Tool calling” is an artificial injection that happens post-training.
- ✘ Too many tools lead to collision.
 - For example, a CRM server exposes both **create_contact** and **update_contact**.
 - The model, told to “add Jane's new phone number,” calls **create_contact** and produces a duplicate record if both descriptions start with “Manage a contact in the CRM...”
- ✘ Token usage skyrockets.
 - 50,000 tokens per tool call for a simple task.

LLMs were never trained for tool calling...



Metaphor: Making an LLM perform tasks with tool calling is like putting Shakespeare through a month-long class in Mandarin and then asking him to write a play in it.

It's just not going to be his best work.

...but LLMs were trained on writing lots and lots of code.

MCP Agent

Traditional tool-calling approach

● Connected Reset End Session

create an event for 6pm tonight for an investor day dinner at la mercerie in soho

```
Tool: create_event
{
  "title": "Investor Day Dinner",
  "description": "Investor day dinner at La Mercerie in Soho",
  "location": "La Mercerie, Soho",
  "date": "2024-01-10T18:00:00"
}
```

Result

```
5,12, / "date": "2024-01-10T18:00:00", / "host_id": "..."
```

Great! I've created your Investor Day Dinner event for tonight at 6pm at La Mercerie in Soho. The event has been successfully created.



```
async () => {
  const today = new Date();
  const eventDate = new
Date(today.getFullYear(), today.getMonth(),
today.getDate(), 18, 0, 0);
  const isoDate = eventDate.toISOString();


  return await codemod.create_event({
    title: "Investor Day Dinner",
    description: "Dinner with investors",
    location: "La Mercerie, Soho",
    date: isoDate
  });
}
```

Now, we need a place to securely execute all of this code at scale.

So, why can't agents just scale on the hyperscalers?

How hyperscalers have scaled...

Microservices were all about vertical scaling, or creating many copies to run the same thing.




From
Monolith

To
Microservices*

** same application,
scaled to millions of users*

...won't work for agents

Microservices rely on being able to run many instances of one application, whereas agents generate *new* instances of applications on *each run*.




From
Monolith



To
Microservices*

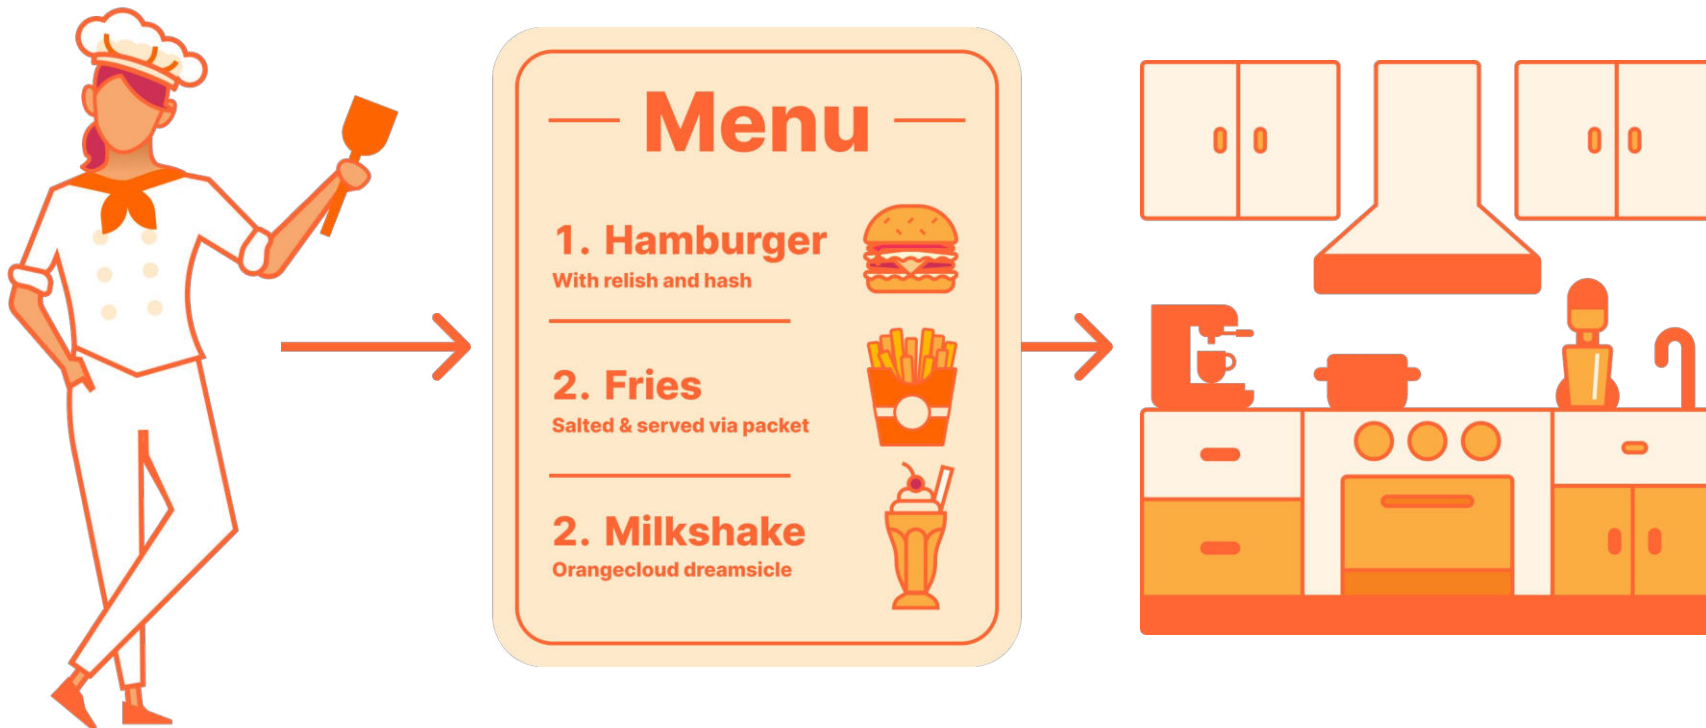
** same application,
scaled to millions of users*



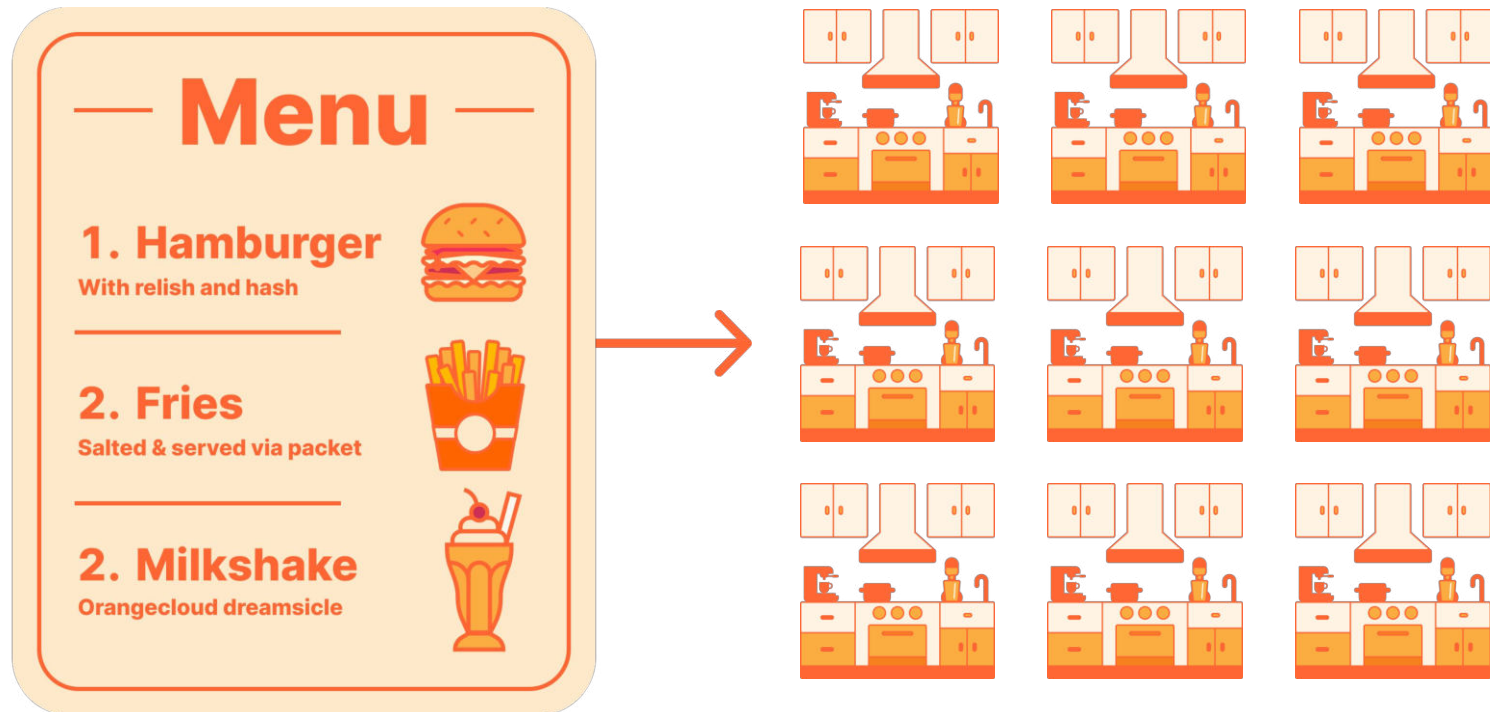
To
Agents**

*** infinite ephemeral applications,
many per user*

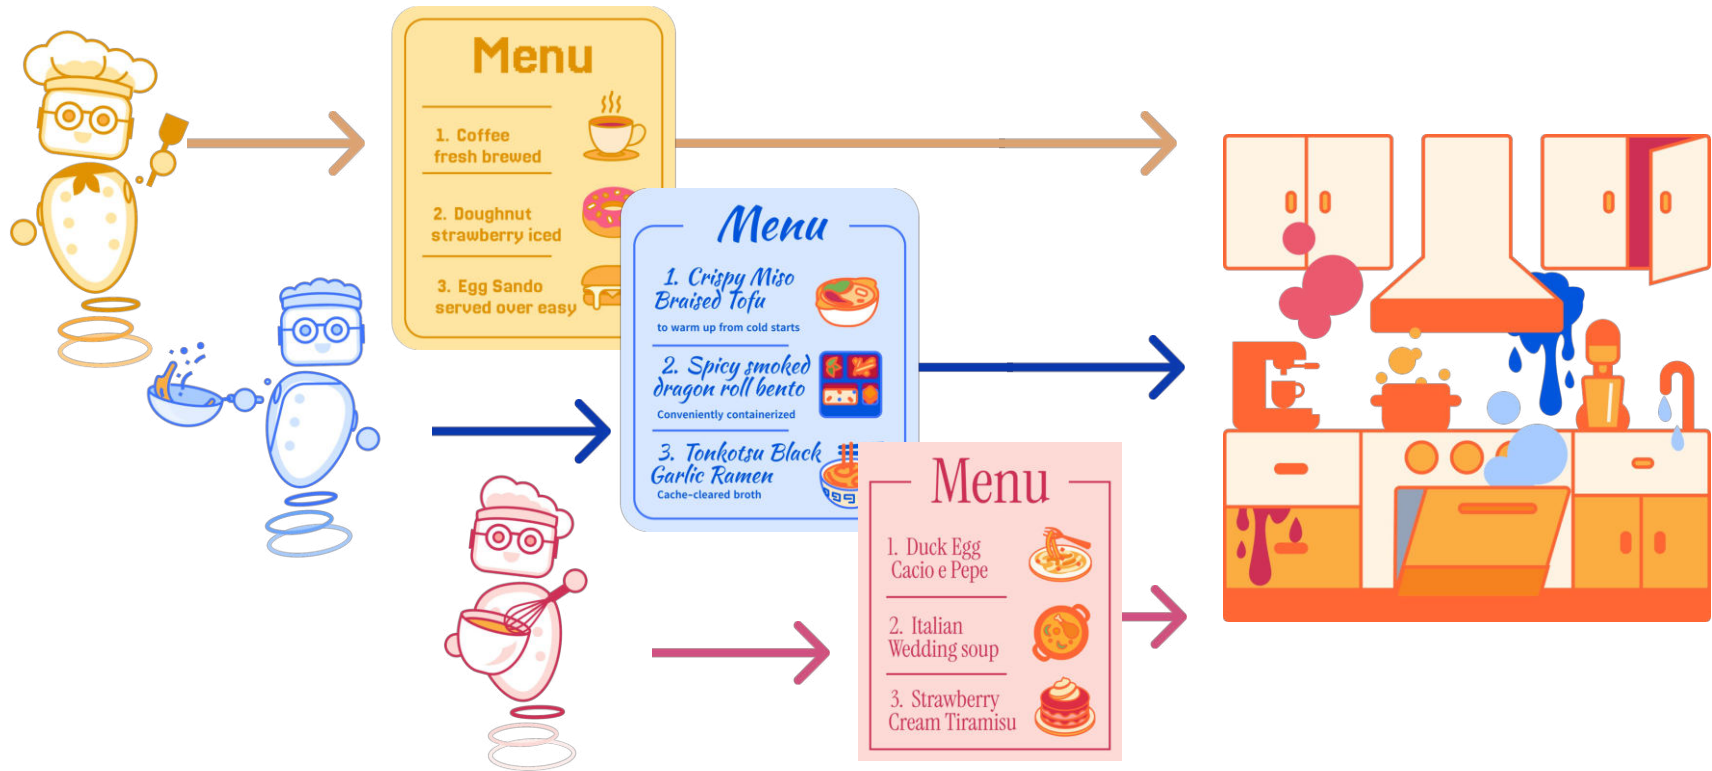
Microservices and hyperscalers: One menu, serving millions per day



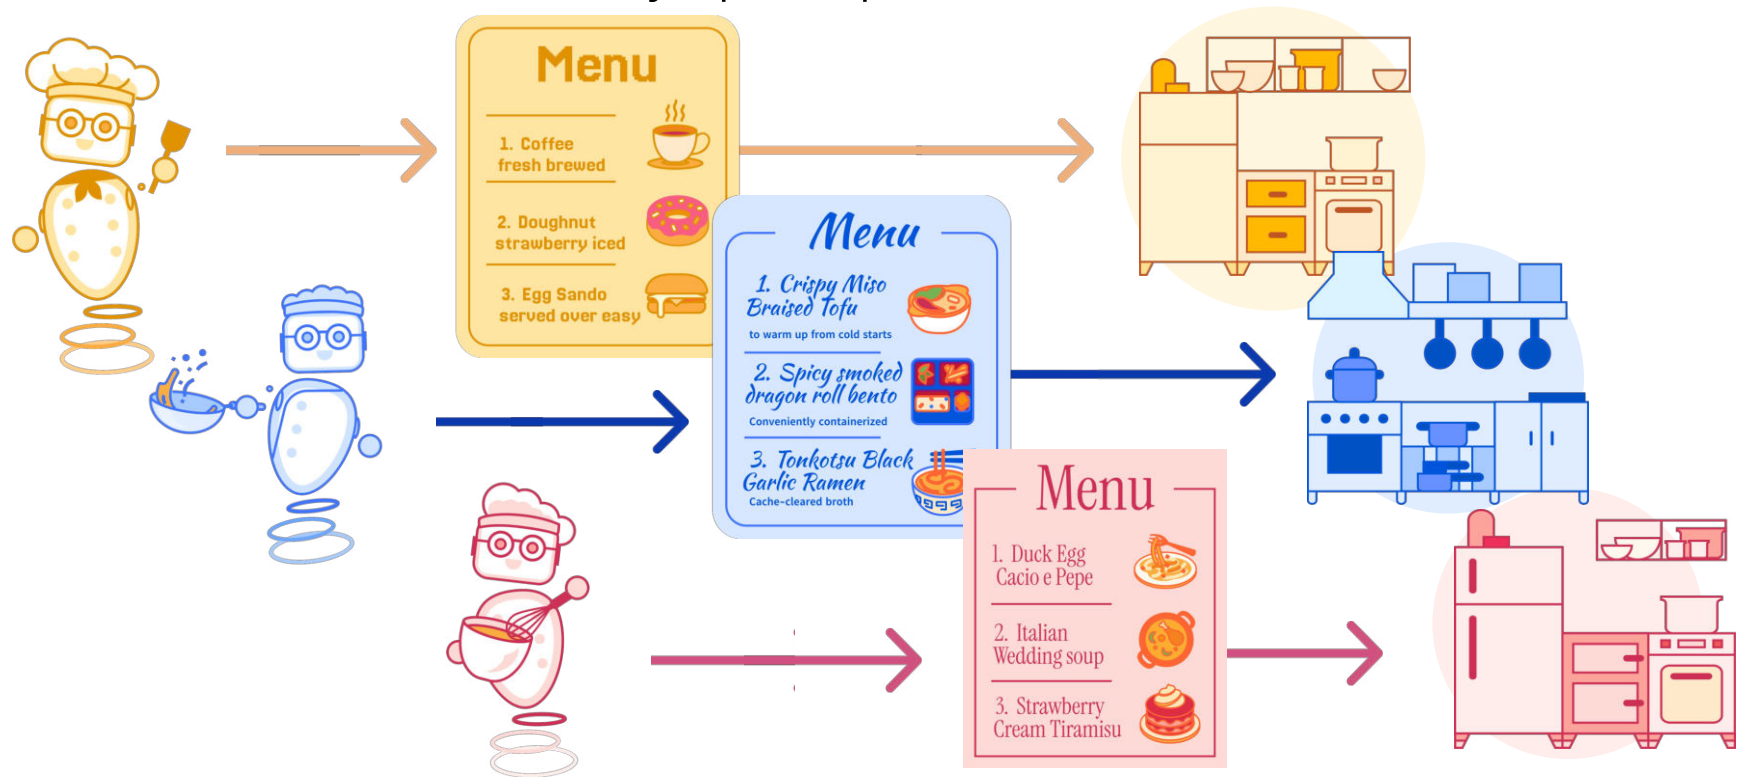
How hyperscalers scale microservices: spin up new franchise kitchens, operationalized around the same menu



The problem for agents with how hyperscalers scale: all compute until today has been designed to execute the same code



What agents need to scale: ephemeral pop-ups, specific to their needs, that can be instantly spun up and down for a task



Why the current cloud approach **won't scale** for agents

United States

100M

Knowledge Workers

1

Agent per worker

~10

Agents per CPU



10M CPUs

Global

1B

Knowledge Workers

10

Agents per worker

~10

Agents per CPU



1B CPUs

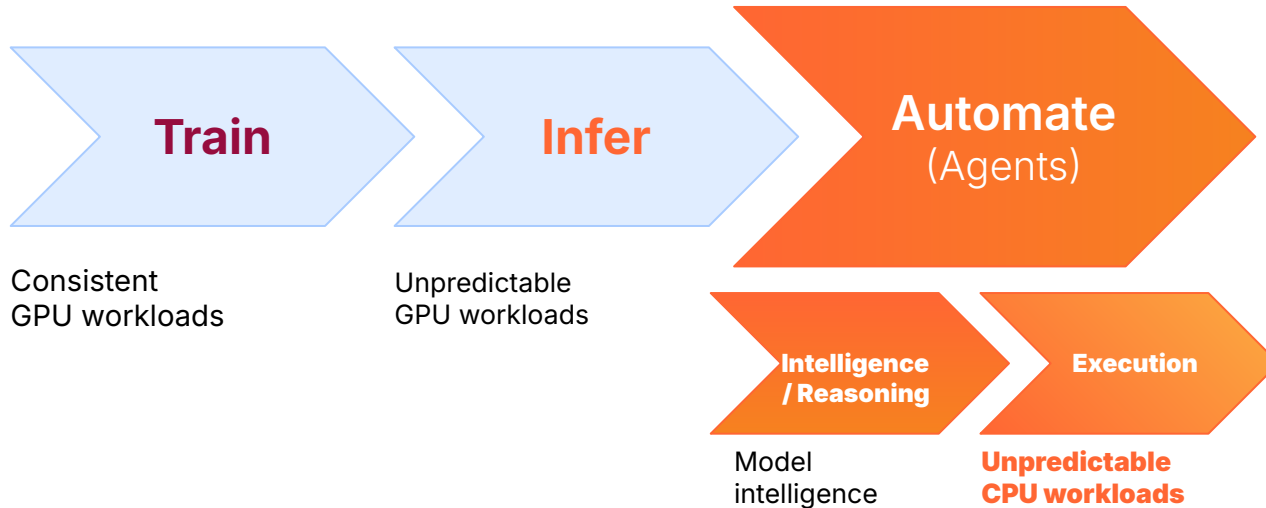
Current global server CPU production:
~35 – 45M per year

~20x current global server CPU production

What's the next bottleneck to AI adoption?

It's the CPU for execution.

— So, how do we scale to billions of agents?

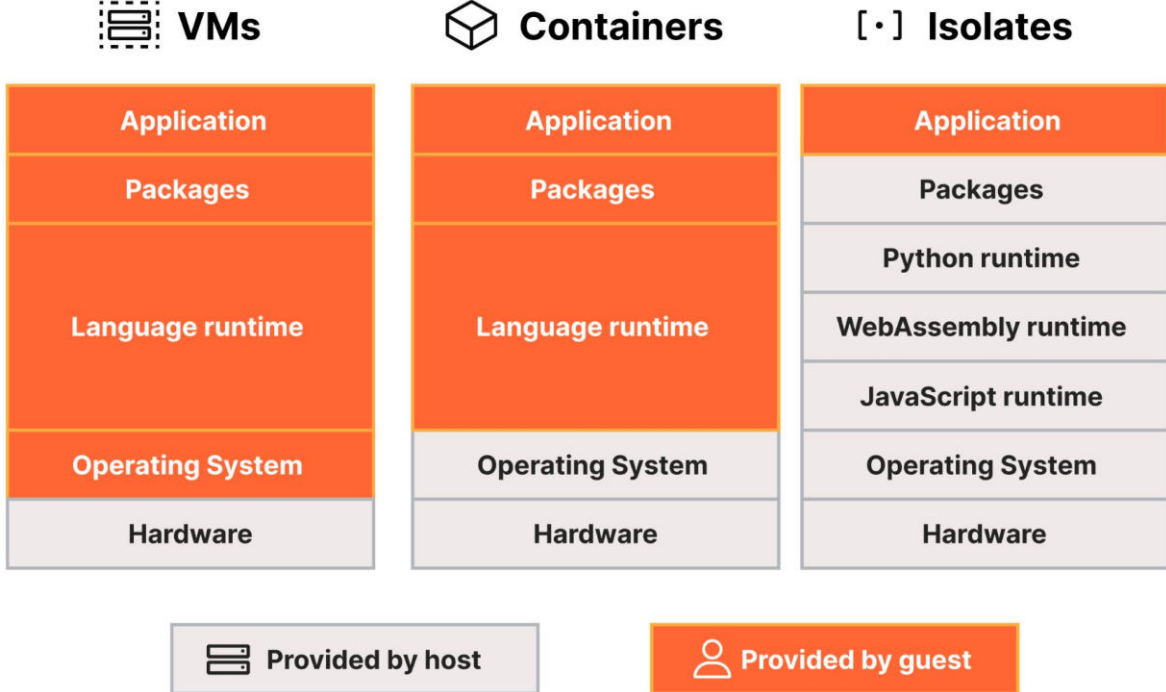


Unpredictable CPU workloads driven by agents is the next bottleneck to AI adoption.


**The cloud that got us here
won't get us to where agents need to go**




Remember that bold, non-consensus decision we made back in 2017?
Isolates were built for this moment!



Because isolates can reshare the runtime against many runs, they can run **up to 100x more efficiently**




 Containers



 Isolate model



User or agent code



Process overhead

What does this mean for customers?

Cloudflare **lowers TCO by 50%+** for agents and the code they write

Serve a single-region web app

1B requests per month

15.2%

lower TCO on AWS

CLOUDFLARE

Workers \$2,301.40

AWS

EC2 Auto Scaling \$1,997.57

VERCEL

Functions \$4,743.84

Deploy 10K agent-generated apps

Every day, for 30 days

63.6%

lower TCO on Cloudflare

CLOUDFLARE

Dynamic Workers \$2,290.95

AWS

Lambda \$6,293.01

VERCEL

Sandbox \$143,138.68

Run 1M agent sessions

Per month

74.7%

lower TCO on Cloudflare

CLOUDFLARE

Durable Objects \$1,444.50

AWS


AgentCore Runtime \$5,718.75

VERCEL

Sandbox \$20,200.60

Source: Cloudflare TCO analysis based on published Cloudflare, AWS, and Vercel pricing.

Andy Jassy was **right about primitives**...



2023 Letter to Shareholders

What matters to builders is having the right tools to keep rapidly improving customer experiences.

The best way we know how to do this is by building **primitive services**. Think of them as discrete, foundational building blocks that builders can weave together in whatever combination they desire. Here's how we described primitives in our 2003 AWS Vision document:

"Primitives are the raw parts or the most foundational-level building blocks for software developers. They're indivisible (if they can be functionally split into two they must) and they do one thing really well. They're meant to be used together rather than as solutions in and of themselves. And, we'll build them for maximum developer flexibility. We won't put a bunch of constraints on primitives to guard against developers hurting themselves. Rather, we'll optimize for developer freedom and innovation."

Of course, this concept of primitives can be applied to more than software development, but they're especially relevant in technology. And, over the last 20 years, primitives have been at the heart of how we've innovated quickly.

Pursuing primitives is not a guarantee of success. There are many you could build, and even more ways to combine them. But, **a good compass is to pick real customer problems you're trying to solve.**

"The best way we know how to do this is by building **primitive services**.

...
Primitives are the raw parts or the most foundational-level building blocks for software developers.

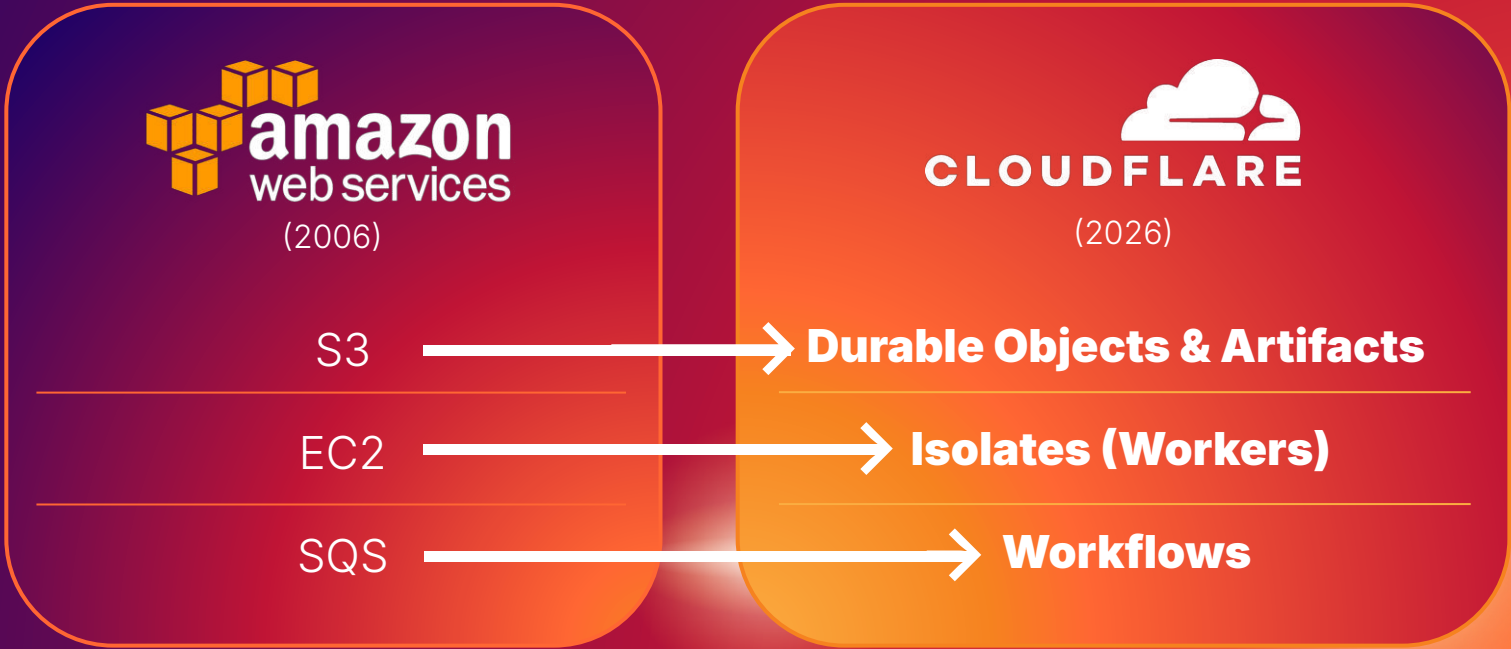
...
And, over the last 20 years, primitives have been at the heart of how we've innovated quickly.

...
Our AWS primitive services were designed from the start to be different.

...
Primitives, done well, rapidly accelerate builders' ability to innovate.

...
a good compass is to pick real customer problems you're trying to solve."

...but the agentic cloud needs new primitives



“

We know what the hyperscalers will look like in 10 years...the same as they do now. We're looking to Cloudflare for what the next generation cloud will look like.”

*– Mark Smith,
VP of Core Tech at Discord*

Internet traffic is changing as a result of agents, driving demand for **agentic cloud primitives to scale**

WHERE:


One Room → Desktop and On-premises → Everywhere → **Anywhere**

HOW:






WHEN:

Batch Deliberate Ubiquitous **Continuous**



Mobile-native businesses
were powered by the cloud


 **airbnb** → 

Uber → 

NETFLIX → 



Agent-native businesses
are powered by Cloudflare

 **Figma** → 

 **Lovable** → 

ANTHROPIC → 

 **OpenAI** → 

In the past month alone...

Figma @figma

wdym of course there's an agent right on the design canvas that's fluent in Figma and native to the way your team works

Lovable @Lovable

Your Lovable apps are now built to be found.

New apps are now server side rendered and discoverable through search engines like Google and AI answer engines like ChatGPT, Claude, and Perplexity. Existing apps get prerendering.

Additionally, access @semrush search intelligence in-app. Chat with your SEO data, understand how you're performing, and improve discoverability.

You can also run SEO reviews on demand, and fix all suggestions in one click.

Claude @claudeai

Live from Code with Claude London: we're launching self-hosted sandboxes (public beta) and MCP tunnels (research preview) in Claude Managed Agents.

Run agents inside your own perimeter, with your security controls applied by default.

402 854 7.5K 2.3M

Claude @claudeai

With self-hosted sandboxes, you can run agents in any environment you control: your own infrastructure, or managed providers like Cloudflare, Daytona, Modal, or Vercel.

OpenAI @OpenAI

Building apps has never been easier.

With Sites, Codex can turn your work, ideas, and plans into an interactive website or app your team can explore, use, and share with a URL.

Rolling out to Business and Enterprise plans, before expanding more broadly.

Cloudflare is the
fastest, most secure, and most cost-effective
place to build, deploy, and scale...

Agents

&

**The code
they write**

8 billion
people → **Multiple** agents
per person → **Infinite**
lines of code

One cloud to power them all

AI@Cloudflare



Sam Rhea

Chief Information Officer

et c0nn3ct c0nn3ct
3ct pr0t3ct pr0t3ct
d 8u1ld 8u1ld 8u1ld 8u1ld 8u1ld


We don't just
build and sell
AI tools and
platforms...





...we are
our own **most**
demanding AI
customer

What is Cloudflare OS?




Equip and enable *every* team member to do more



One-Click AI Agent Workspace

Full agent harness running in a Cloudflare container available to every employee seconds after they start work. Build apps, outputs, and manage agents.




Centrally Managed Context Layer

Curate the jobs-to-be-done across the organization, and our institutional knowledge, into material that any team member can use in their workspace.



Secure System of Record Access

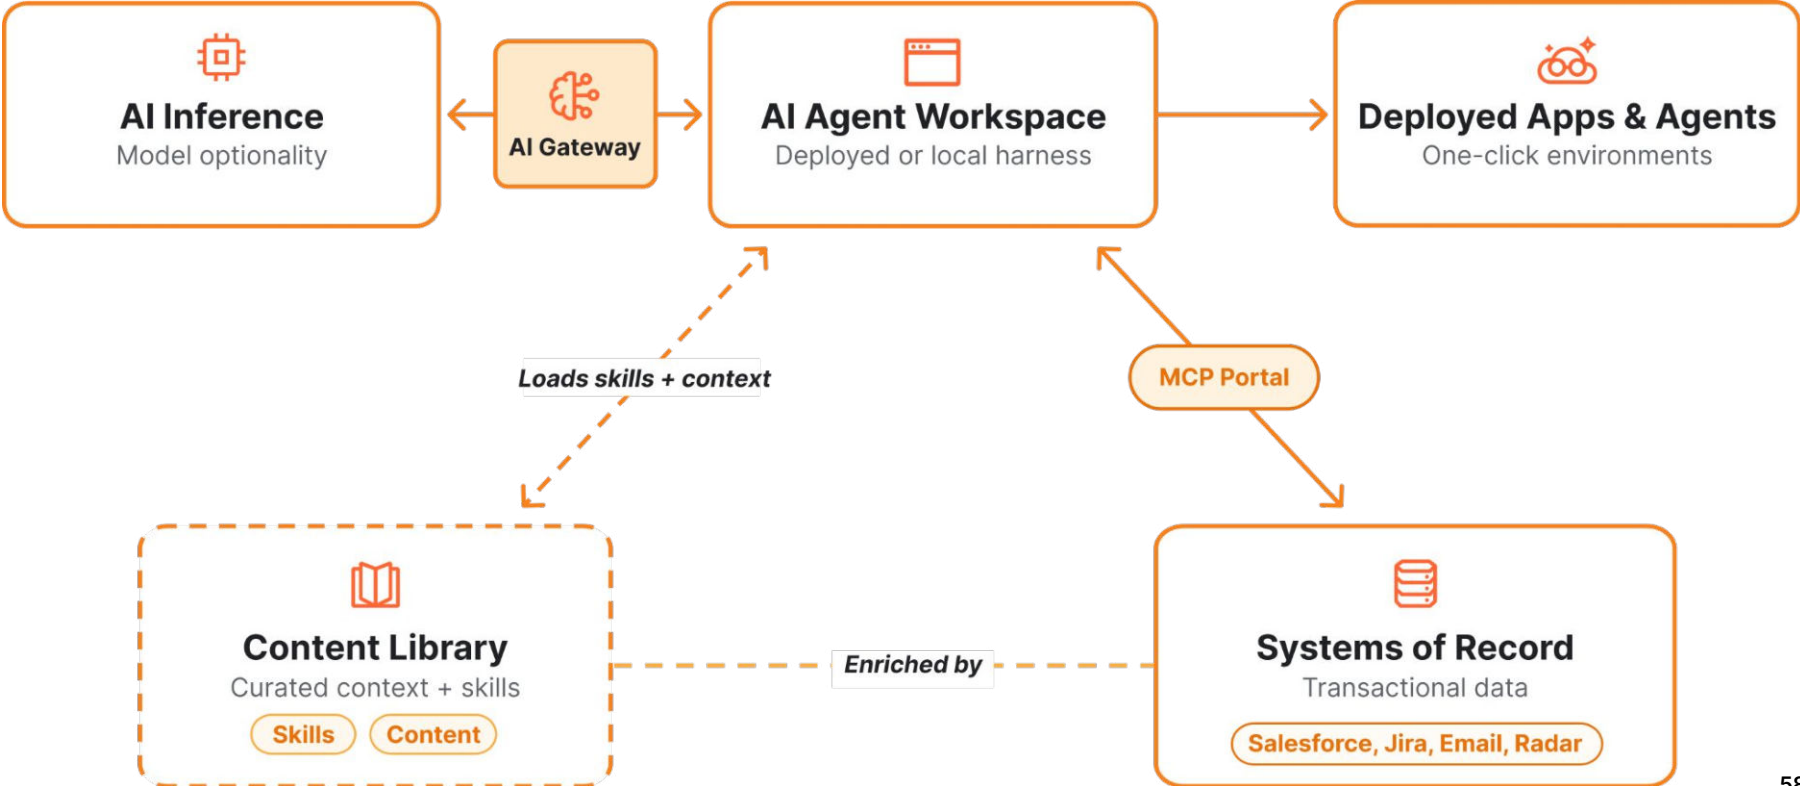
Provide access to critical data and systems of record across the organization while tightly scoping what a user can access and what they can do.



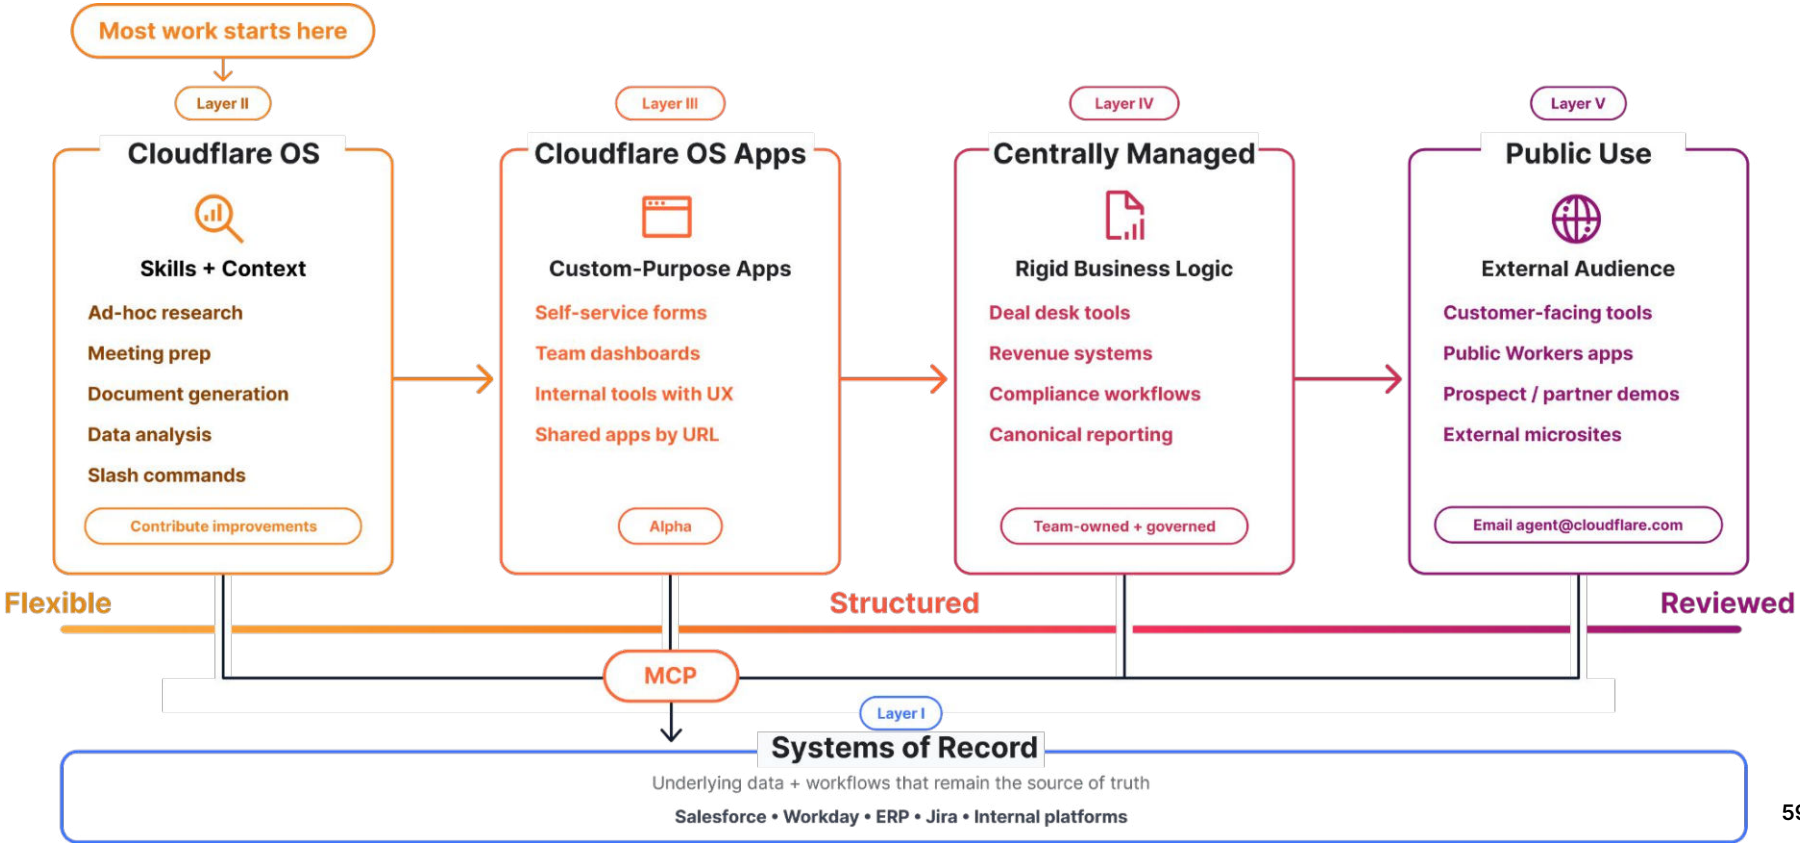
Model Optionality by Use Case and Role

Default to the right model by task or role to control cost without slowing down productivity gains.

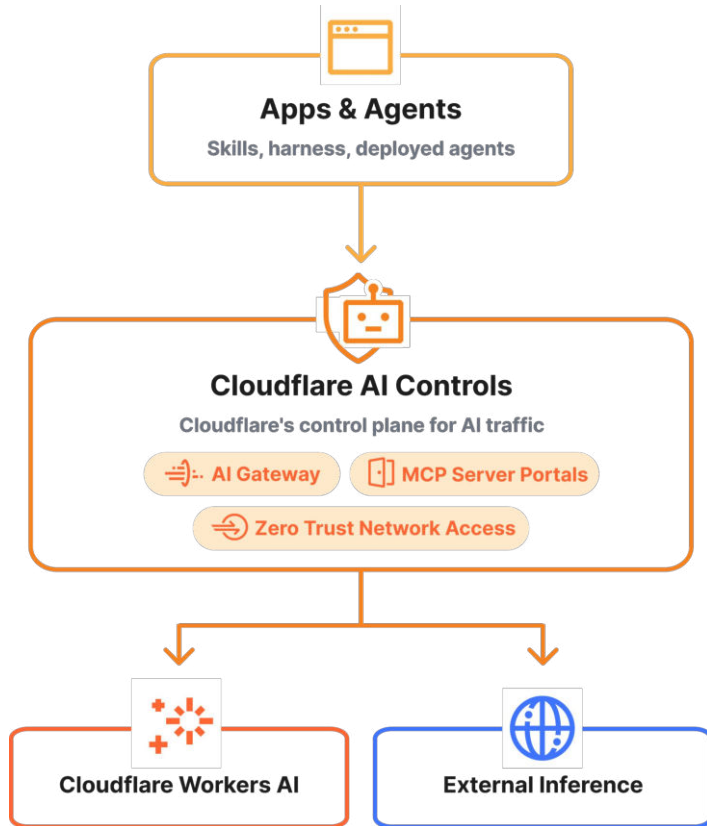
Cloudflare's agentic workspace



A new framework for where and how team members build



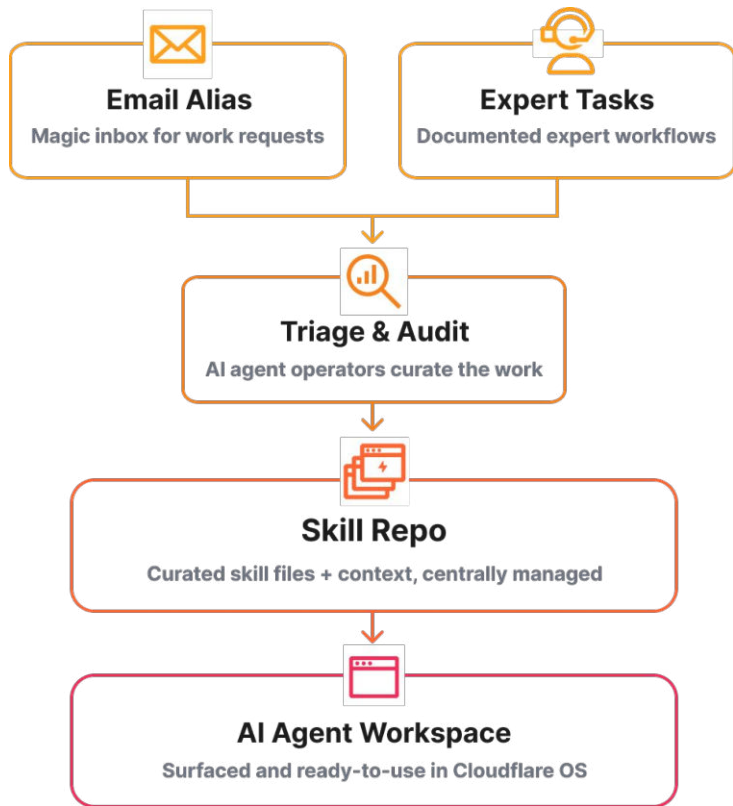
Scaling productivity without scaling cost



Cost-efficient AI usage for every employee:

- ✓ **Model optionality**
Control model usage without lock-in risk
- ✓ **Dynamic routing**
Map employee roles and skill files to specific models
- ✓ **Infire**
Faster, more cost-efficient inference
- ✓ **Caching**
Save token spend on at-scale usage

Investing in a central context layer



Defining the jobs-to-be-done as centrally managed skill files:

- 1 Tap experts to document workflows and create skill files.
- 2 Create a "magic inbox" email alias where everyone else can send work.
- 3 Staff the inbox with AI agent operators who run it as a service.
- 4 Collect and curate the use cases, convert to skill files.
- 5 Everything is surfaced and ready-to-use in Cloudflare OS.

Cloudflare OS in *Action*





A comprehensive skill library

Every team member

Workspaces

Outputs

Scheduled Tasks

Context

Skills

Search...

+ New workspace

> ☆ FAVORITES 0

> ⌚ RECENT 10

 v Show 28 more

> 📁 SCHEDULED TASKS 1



What would you like to do?

Start a new workspace for your task

Ask Seal to research, write, analyze, or create something...

+

GLM 4.7 Flash ▾





Builders



Sellers



Code review lets us ship even faster, safely, safely

Builders

Trust / Merge requests / I2123

Open Draft: DS-16240: Migrate Apollo & Oxy to lassen protocol v2 adepke/DS-16240-migrate-la... into master

Overview 4 Commits 1 Pipelines 1 Changes 11

3 open threads ^ v ⋮

Add a to-do item

Activity

All activity v ↑

Andrew Depke assigned to @adepke 1 month ago

AI Code Reviewer @gsa_claude 1 month ago

Developer 🗨️ ↶ ⋮

OpenCode Code Review for 681241de

Approved with Comments — Clean feature-gate plumbing following existing patterns. The sentinel-value hack in clear_child_log_id() and the unused lassen_v2 field on SessionLogChildInfo are the main design concerns — both acknowledged as draft-quality and straightforward to fix before leaving Draft.

Code Quality: ⚠️ 4 warnings

Sentinel-value pattern for v2 child-clearing is fragile; SessionLogChildInfo.lassen_v2 is populated but never consumed; silent fallback masks logic errors.

Recommendation: ⚠️ Major Comments

- Sentinel value [0u8;16] in clear_child_log_id() v2 path collides with the zeroed parent_id in build_lassen_log_message and makes get_child_log_id() return a fake ID — replace with a dedicated AtomicBool child_cleared field as the code comment itself suggests
- SessionLogChildInfo.lassen_v2 is set in get_session_log_child_info() but never read by any consumer — either wire it through to Gateway's request proto or remove it
- unwrap_or_else(LassenMessageId::new_root) silently generates a fresh root ID when root_message_id is None in v2 mode — for TCP/UDP this masks a logic error; use invariant_violated! with safe fallback instead
- Import reordering in ip_tunnel_context.rs and app_ctx.rs is cosmetic noise — consider a separate commit

Assignee Edit

Andrew Depke

0 Reviewers Edit

Requires 1 approval Assign

None - assign yourself

Labels Edit

None

Milestone Edit

None

Time tracking ⌚ +

No estimate or time spent

2 Participants





Agents fix problems so engineers can focus on building Builders

Send feedback

Edit

Sandbox Files Skills Tasks Access MCP

Workspaces

Outputs

Scheduled Tasks

Context

Skills

Search...

+ New workspace

FAVORITES 1

RECENT 10

Show 6 more

TODAY

2:53 PM

fix the issues found in this MR
https://gitlab.cfdata.org/cloudflare/cto/cto-agent/-/merge_requests/916

2:55 PM

> 2 code executions

- Executed code `fetch MR 916 details, discussions, and notes to identify issues` 4 calls · 4.5s
- Executed code `get MR basics like title branch source/target and state`

> 2 tool calls

- proxy_fetch GET `gitlab-access.cfdata.org/api/v4/projects/cLoudfLare%2Fcto...`
- proxy_fetch GET `gitlab-access.cfdata.org/api/v4/projects/cLoudfLare%2Fcto...`

I see the MR — Aaran's Sentry sourcemap upload CI change. The OpenCode review flagged 2 blocking issues + 2 suggestions, and the pipeline failed. Let me get the rest of the context and check the pipeline before dispatching.

> 3 tool calls

- proxy_fetch GET `gitlab-access.cfdata.org/api/v4/projects/cLoudfLare%2Fcto...`
- proxy_fetch GET `gitlab-access.cfdata.org/api/v4/projects/cLoudfLare%2Fcto...`
- provision_sandbox `sb-e95e5cc74980`

Thinking...

Send a message...

Seal model Claude 4.7 Opus

Sandbox model Claude 4.7 Opus

Stop

Send



BDRs research prospects at scale for outbound work

Sellers

← Back to search

Research will continue in the background



Target

target.com

Q Normal

- Gathering basics**
Quick company info & site analysis
- Searching the web**
News, company intel & tech stack
- Finding content**
Blog posts & case studies
- Analyzing**
Strategy, ICP scoring & opportunities
- Building report**
Saving the final report

All Reports

Bulk

Account Tiers



Every seller starts their day with a prescriptive plan

Sellers

```

tool_7n6K2Z6d_salesforce_... SELECT Customer_Segment__c, Customer_Health_Score__c, Acc
tool_7n6K2Z6d_salesforce_... SELECT Name, Amount, ACV__c, CF1_ACV__c, CloseDate, Stage
tool_7n6K2Z6d_salesforce_... SELECT Name, Amount, ACV__c, StageName, CloseDate FROM Op
tool_7n6K2Z6d_salesforce_... SELECT Name, StageName, CloseDate FROM Opportunity WHERE
tool_7n6K2Z6d_salesforce_... SELECT CaseNumber, Subject, Priority, Status, CreatedDate
tool_7n6K2Z6d_salesforce_... SELECT CaseNumber, Subject, Status, ClosedDate, JIRA__c F
tasks:update sfdc-pull +2 more · status: completed
tool_7n6K2Z6d_wiki_mcp_se... text ~ "Target Corporation" AND type = "page" or.. limit:..
fs:list cloudflare/context
fs:write memory/target-call-prep-data.md
fs:list cloudflare/context/facts
fs:list cloudflare/context/proof
fs:list cloudflare/context/mappings
fs:list cloudflare/context/Verticals
fs:read cloudflare/context/facts/competitive.json
fs:read cloudflare/context/proof/case-studies.json
fs:read cloudflare/context/mappings/product-mapping.md

```

I have enough to draft. Writing the prep doc now.

```

> tasks:update
update research-target +2 more
status: completed

```

Thinking...

Send a message...

Claude 4.7 Opus

Stop

Send

No documents yet
 Ask the agent to draft something, or click + to start

Files

No documents yet
 Ask the agent to draft a document, or click + above to start.
 Documents appear here from this workspace.



Generate custom slide decks for each conversation

Sellers

Connected Disconnect

Session

Changes

Cloudflare Developer Platform for Retail

I need a slide deck for a retailer who operates physical locations and has a large digital presence about how they should adopt Cloudflare's developer platform

I need to ask you which visual style you'd like before I start creating slides.

Questions 1 answered

Which visual style would you like for this retail developer platform deck?

Cloudflare Corporate Template (Recommended)

Ask anything...



Build

Claude Opus 4.6

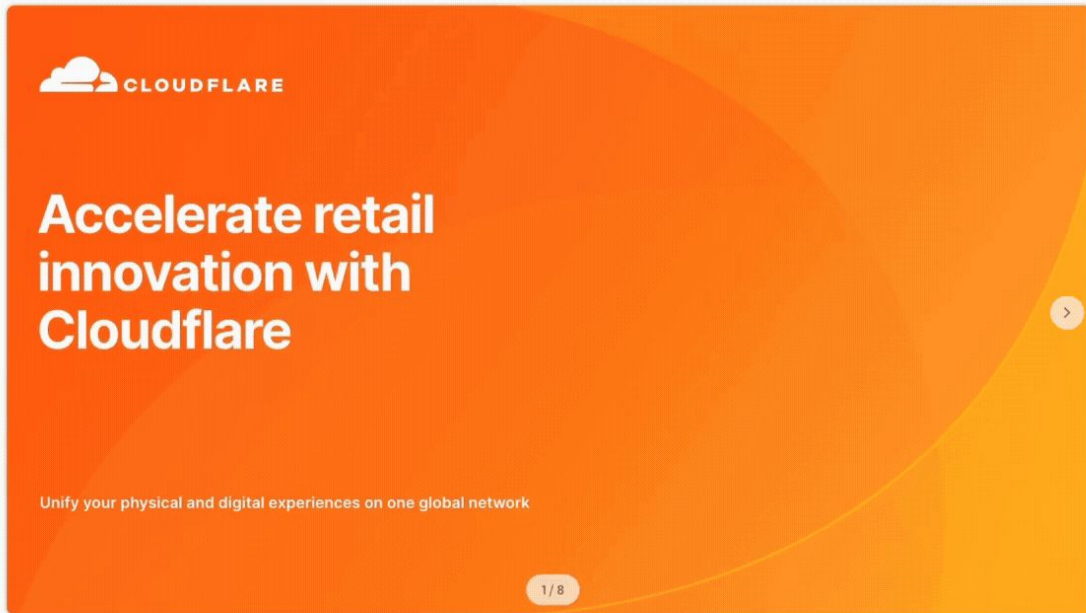
Default

Images

Slides

8 slides

- 1 1-cover
- 2 2-retai...
- 3 3-platf...
- 4 4-solu...



1 / 8



Spend event planning time with customers instead

Dashboard

Sellers health, regional coverage, and staffing signals

Add event

Import

BDR Align

Execs

Design

Lead Follow-up

Partner

Developer

Regions

Settings

TOTAL EVENTS

483

BOOKED

142

BDR REQUIRED

197

EXECUTIVES ONSITE

0

DESIGN REQUIRED

0

IN EVALUATION

105

IN PROGRESS

0

NEXT 30 DAYS

0

COMPLETED

2


Regional Coverage

Events by global region



Strategic Pillars

Both event pillar taxonomies are supported



Upcoming Events

Next booked or planned dates across regions

No upcoming events yet.



Sellers create their own solution architecture proposals

Sellers

TODAY

2:00 PM

Using /solution-brief

No documents yet

Ask the agent to draft something, or click + to start

Files

Workspaces

Outputs

Scheduled Tasks

Context

Skills

Search...

+ New workspace

FAVORITES 0

RECENT 10

Show 20 more

SCHEDULED TASKS 1

Task: sam's call prep

for a pharma customer replacing their traditional vpn with Cloudflare Zero Trust

2:01 PM

Executed code gather Cloudflare context for pharma VPN ... 6 calls · 216ms

- fs:read cloudflare/context/facts/competitive.json 33ms
- fs:read cloudflare/context/facts/network.json 34ms
- fs:read cloudflare/context/proof/case-studies.json 36ms
- fs:read cloudflare/context/references/writing-guidelines.md 36ms
- fs:read cloudflare/context/facts/analysts.json 36ms
- fs:read cloudflare/context/voice/writing-guidelines.md 41ms

7 tool calls

- skills:load solution-brief
- skills:load solution-brief
- fs:read cloudflare/context/facts/network.json
- fs:read cloudflare/context/facts/analysts.json
- fs:read cloudflare/context/facts/competitive.json
- fs:read cloudflare/context/proof/case-studies.json
- fs:read cloudflare/context/voice/writing-guidelines.md

Thinking...

Send a message...

Kimi K2.6

Stop

Send

No documents yet
Ask the agent to draft a document, or click + above to start.
Documents appear here from this workspace.



Sellers position Cloudflare by use case and region

Sellers

- Workspaces
- Outputs
- Scheduled Tasks
- Context
- Skills
- Search...
- + New workspace
- > ☆ FAVORITES 0
- > ⌚ RECENT 10
 - ▼ Show 27 more
- > 📁 SCHEDULED TASKS 1

TODAY

2:31 PM

Using /case-study-selector

competing against Zscaler for a SASE deployment at an asian manufacturing company

2:32 PM

> 3 tool calls

- skills:load case-study-selector
- fs:read cloudflare/context/proof/case-studies.json
- fs:write local/case-studies-sase-zscaler-asian-mfg.md

Thinking...

Send a message...

Kimi K2.6

Stop Send

case-studies-sase-zscaler-asian-mfg.md

Markdown - 5.8 KB - just now

Files

Source Page Branded Save to Outputs

Case Study Selector — SASE vs. Zscaler | Asian Manufacturing


Notes

Scenario: Competing against Zscaler for a SASE deployment at an Asian manufacturing company. Framing: SE / AE discovery call with SASE / Zero Trust focus.

Primary: Canva — APAC Zero Trust at Scale

ATTRIBUTE	DETAIL
Industry	Technology / SaaS
Region	APAC
Products	Access, WARP, Gateway, WAF, DDoS, Bot Management, RBI, Tunnel

And, we're just getting started



Act IV



Stephanie Cohen

Chief Strategy Officer

Cloudflare was **built for this moment**



**The business model
of the Internet is changing**



**Cloudflare is the only cloud
built for the Agentic Internet**



**Every agentic interaction
on the Internet is a chance
for commerce**

...and we're already seeing results



Cloudflare

Platform of choice for the agentic web

Orange is the new cloud

We are already seeing results

Media sector proof points

37%+ revenue growth

in Q1'26 vs. Q1'25

1.3x higher revenue per customer

than company average

60%+ expansion rate

for retained customers in Q1'26

~117% DNR

in Q1'26



**Win new
customers**



**Grow existing
customers**




**Increase
customer
stickiness**

The business model of the Internet is **changing**



AI is a platform shift, and the adoption is unprecedented


Years to reach 1 billion users globally



2x+ faster


















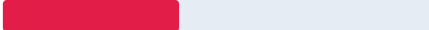

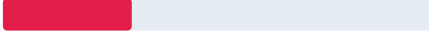
adoption of Gen AI than any other preceding platform shift

AI adoption is quickly transforming where we spend our time



Note: Illustrative / directional framing, based on the sources below — not a single measured series. 2000s/2010s shown as ~100% open web pre-social / pre-AI; 2015-22 split approximates attention shifting to social. Measured anchor: Pew Research Center, AI-summar click-through study (U.S. browsing data, Mar 2025). Directional context: DataReportal Digital 2023-2024; SparkToro & Similarweb Zero-Click Search Studies; data.ai.

AI is **breaking the business model** of the Internet


Publication	Peak Performance	Jan 2026	Traffic Decline
 digitaltrends	Mar 2024 8,530,891	264,861	 -97%
 ZDNET	Feb 2024 7,610,480	768,792	 -90%
 TheVerge	Feb 2024 5,322,037	790,002	 -85%
 How-To Geek	Feb 2024 1,974,331	293,898	 -85%
 techradar	Jul 2024 15,577,298	4,045,783	 -74%
 WIRED	Nov 2024 7,754,067	2,976,994	 -62%
 tom's guide	Jul 2024 16,013,790	7,986,572	 -50%
 CNET	Nov 2024 20,294,300	10,655,803	 -47%
 PC MAG	Jul 2025 12,667,236	7,449,728	 -41%
 Mashable	May 2024 16,114,803	11,331,018	 -30%

It's not **"just"** the media industry



We are seeing **multiple industries** impacted

Absolute human traffic by Industry (June 2025 vs. April 2026 comparison)



**35 - 40%
DECLINES**



Source: Cloudflare Radar — absolute human (non-bot) request volume by industry, June 2025 vs. April 2026.

And, the business of the Internet **needs saving**

6+ billion

people connect to the Internet, a number that's doubled in 11 years

\$16+ trillion

value of the digital economy, 15% of global GDP

59%

of people use the Internet to stay connected with friends and family globally

61%

of users cite information access as the top reason for Internet use




Mission:

Help build a better Internet



Cloudflare is the **only cloud** **built for** the Agentic Internet





Supply

Transparency • Control • Optimize • Charge



Bot Management

Filter and manage automated traffic



Web Bot Auth

Authenticate legitimate bots



Content optimization

Transform content to Markdown for Agents



Technical Capabilities

Reliable scarcity

Analytics

Value discovery

Neutrality / a level playing field

Streamlined workflows

A payment system



Demand

Optimize • Behave • Access • Discover



Infrastructure

AI Gateway, Workers, & Workers AI



Developer Tools

Agents SDK, MCP, AI Search



Visualization

Browser Rendering for Agents

Cloudflare is positioned to provide
the “control plane” for the Agentic Internet



Supply


The Internet runs on Cloudflare

Breadth




20%+ of the Internet runs on Cloudflare

Depth



36% of the top 10,000 sites run on Cloudflare

Importance



42% of the Fortune 500

Uniqueness



The world's largest developer communities run on Cloudflare

The economy runs on Cloudflare



MEDIA



SONY MUSIC



indeed

reddit

stack
overflow

AXIOS



SOFTWARE & SAAS

Firefox

Discord

zendesk

ORACLE

IBM

BROADCOM



RETAIL & E-COMMERCE

Carrefour

shopify

Delivery Hero

Uber

DOORDASH

FOSSIL



FINANCE & FINTECH

SoFi

BANCO DO BRASIL

Investec

SAISON
CARD

Bank of Cyprus

workday



GAMING & GAMBLING

CCP

playtech
SOURCE OF SUCCESS

mod.io

SYBO

kaizen
GAMING

Caliente
INTERACTIVE



HEALTHCARE

labcorp

Doctolib

Flo

VERTO

Fullscript

Roche



Demand


Leading AI companies already run on Cloudflare too

Top 50 generative AI Web products, by unique monthly visits



Developers increasingly build their agents on Cloudflare

Agents SDK Downloads



42x
growth in monthly
downloads

14M+
downloads since
launch

Source: npm public download statistics for the agents package (Cloudflare Agents SDK; github.com/cloudflare/agents) from March 17, 2025 through June 2, 2026.

The **foundational beliefs** that built our platform



Secure by default

A property of where your code runs, not a product you attach to it



Fast by default

A consequence of the platform being everywhere, not a region you pick



Scalable by default

Handled by the platform, not capacity you provision and manage



Compliant by default

Data stays where it needs to stay, without you rebuilding for every jurisdiction



Verified by default


Trust on every request, not assumed because someone is inside a network



Open by default

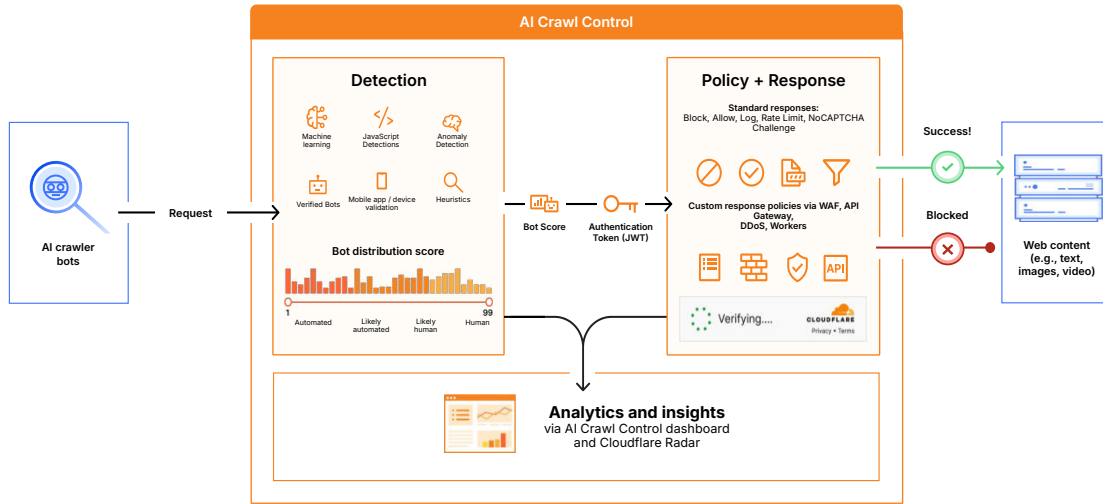
Your data moves freely, not locked in by egress fees or proprietary barriers

Technical capabilities required for a functioning market



We are the best at identifying, classifying, and controlling bots

Our competitive advantage is visibility. The more of the Internet we see, the better we identify, classify, and stop unwanted or malicious bots — before they ever reach our customers' origin.



Registered bots
500+ registered bots identified

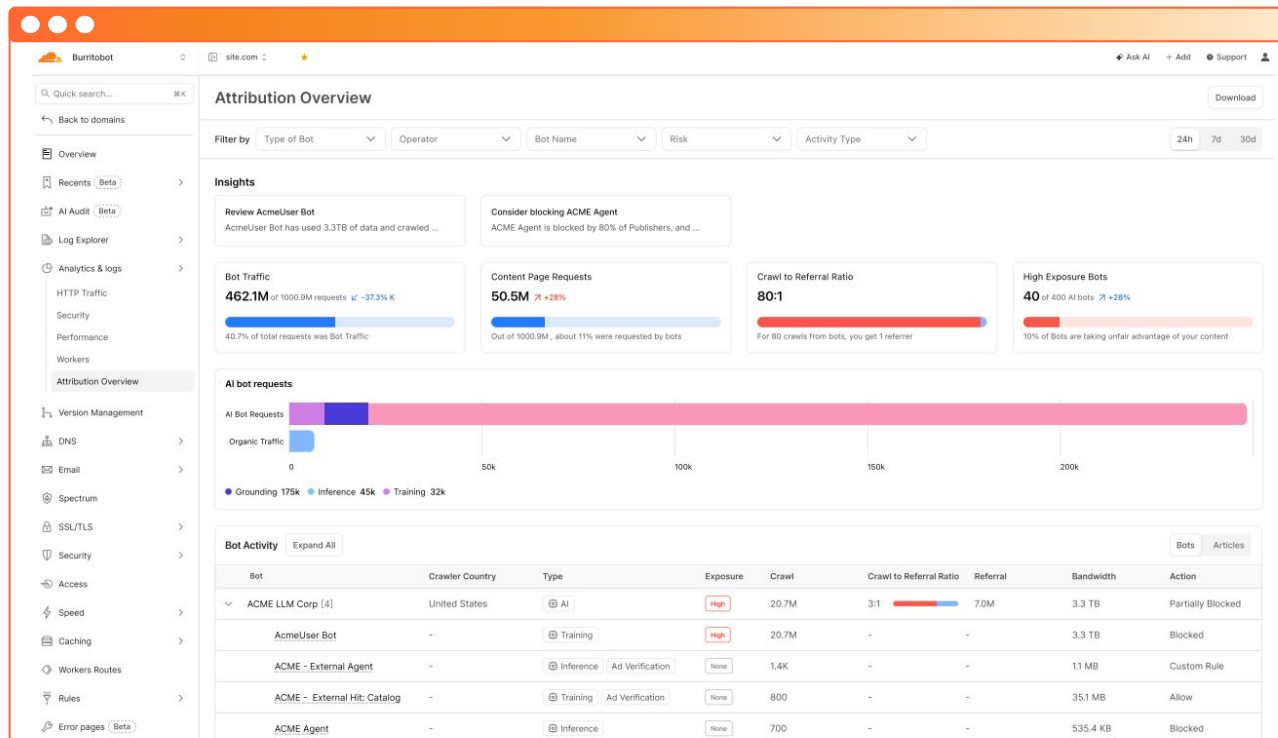
Requests processed
Hundreds of billions of bot requests processed daily

See it first
Identify new bot signatures across 20%+ of the web instantly

Learn faster
ML models trained on billions of daily interactions

Block smarter
Stop bad bots before they scrape, spam, or abuse


Transformation of data to insights is key for business development



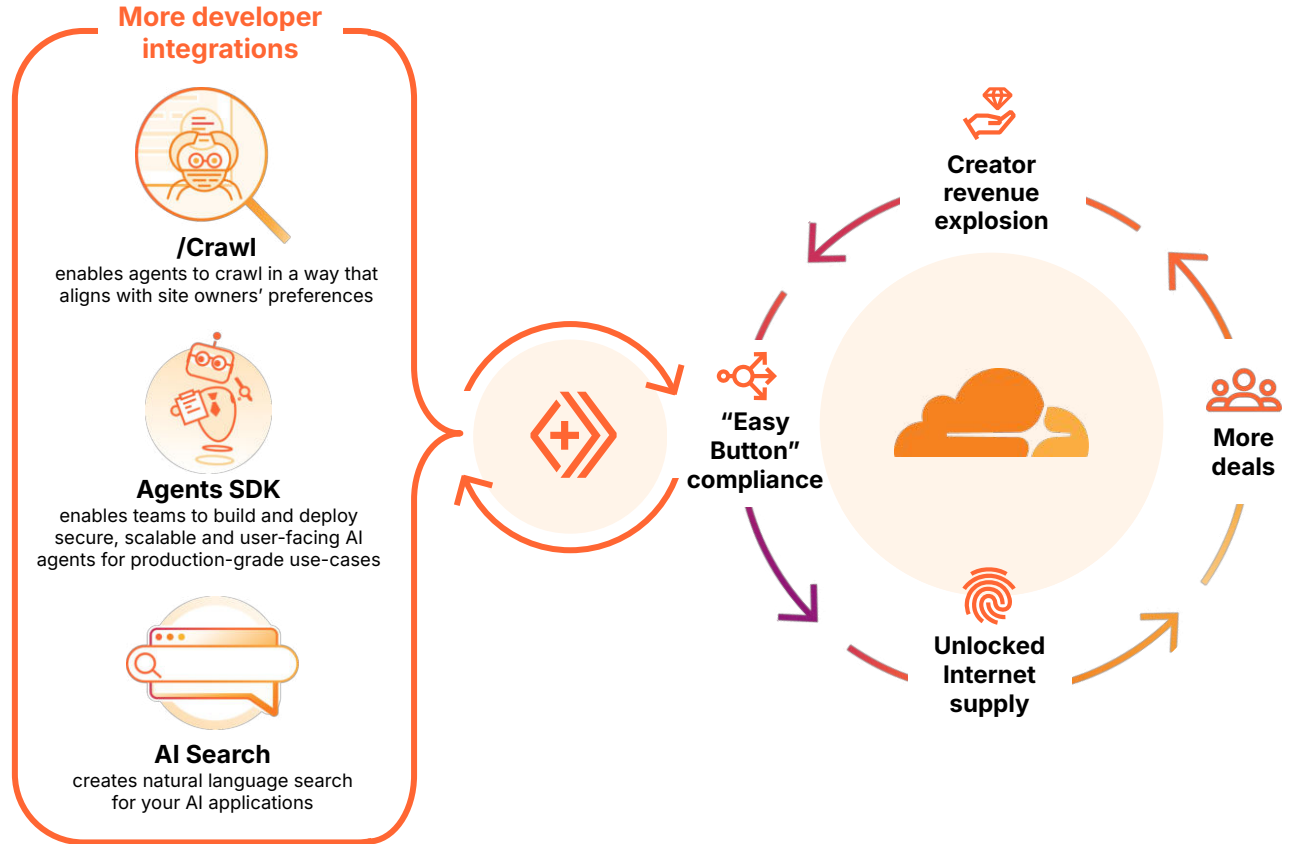
Bot Management began as a threat intelligence product for security but is now **also being** used for market intelligence to drive marketing and partnerships decisions.

First-party signals identify **highest-quality** content

Cloudflare's network processes trillions of requests daily, generating unique signals that help partners **identify and surface the most relevant information**.



“Easy Button” for well-behaving crawlers & agents



Building agents on Cloudflare means your code runs milliseconds away from the user.

Cloudflare is the only network where well-behaving crawlers and edge-deployed agents are the default.

Cloudflare **does not compete** with publishers or AI labs


As a neutral third-party, Cloudflare can **serve both sides**...

Neutrality transforms Cloudflare from an infrastructure vendor into a market maker for content and agentic access:

- ✓ **No frontier LLM**
- ✓ **No publisher competition**
- ✓ **Trusted by both sides**

Unique page views by AI crawlers

Multipliers show less access vs. #1 crawler



...and our unique position also enables us to provide a **level playing field** across frontier LLMs.

#6: FUNCTIONING MARKETS REQUIRE A PAYMENT SYSTEM



LARGEST PAYMENT NETWORK

900 million

Average transactions per day

<100,000 transactions per second



CLOUDFLARE

30+ trillion

Average transactions per day

500+ million transactions per second

Cloudflare is the **bridge** with the traditional payments ecosystem

Web Bot Auth enables merchants to cryptographically verify agent identity before processing transactions.

PAYMENT NETWORK ADOPTERS



Trusted Agent Verification

Deliver secure agentic commerce experiences via the Web Bot Auth standard

cloudflare/web-bot-auth

The Open Standard & Foundational Protocol Base Layer


Unlike API keys that authenticate systems, Web Bot Auth cryptographically binds identity to specific transaction flows, enabling fine-grained authorization policies for autonomous commerce.

Every agentic interaction on the Internet is a **chance for commerce**



The agentic surge is real

AI Agent Requests (in billions per day)



1,700%+

INCREASE

in daily AI agent requests
on Cloudflare's network
over one year

Machine payments will grow with machine traffic

Cloudflare founded the x402 Foundation with **Coinbase and Stripe**.

x402 can become the standard to connect machine buyers and sellers in the request path.

2+ BILLION
402 responses served per day
on Cloudflare's network



AI agent requests resource

API, content, or tool
via MCP



Cloudflare returns

402 + price

Payment required



Agent signs & retries with 402

Cryptographic commitment



Edge verifies & serves

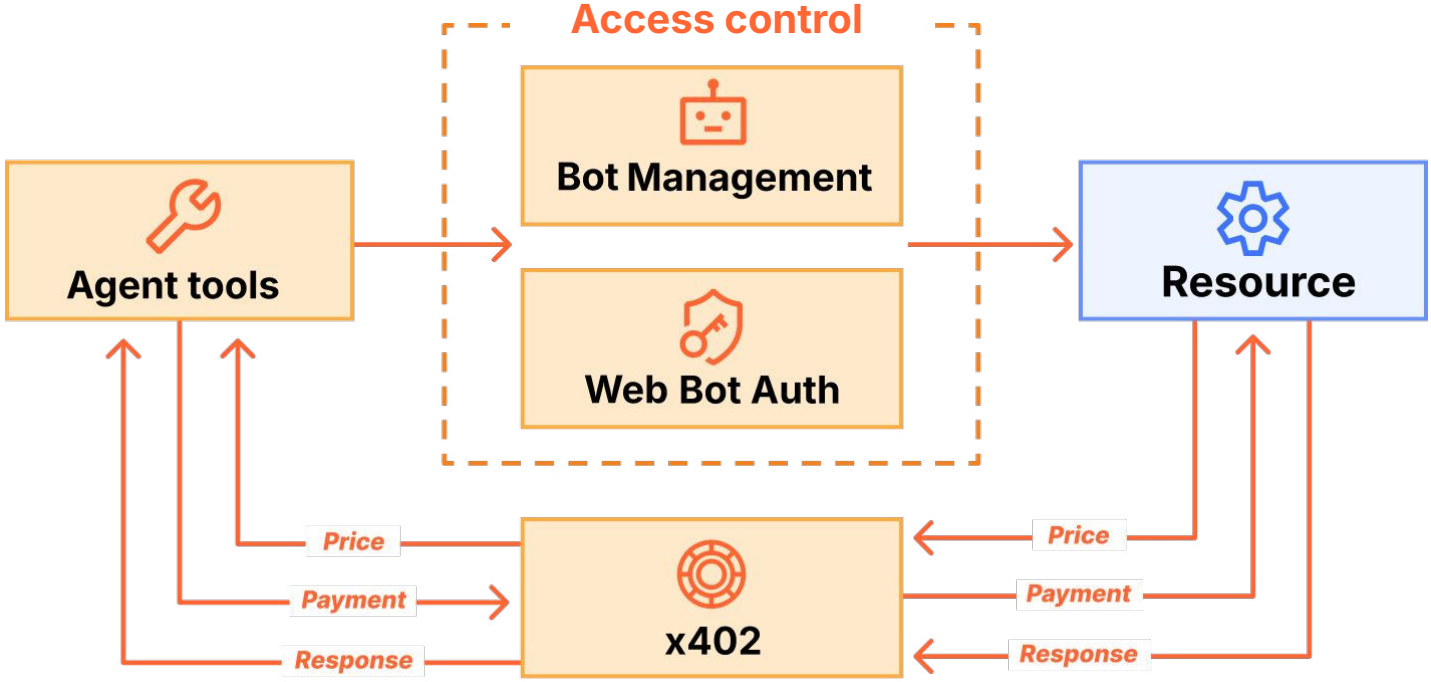
Payment checked at
network speed

Success: Payment is verified at the edge before the agent ever receives HTTP 200 access.

Source: Cloudflare network data, May 2026.



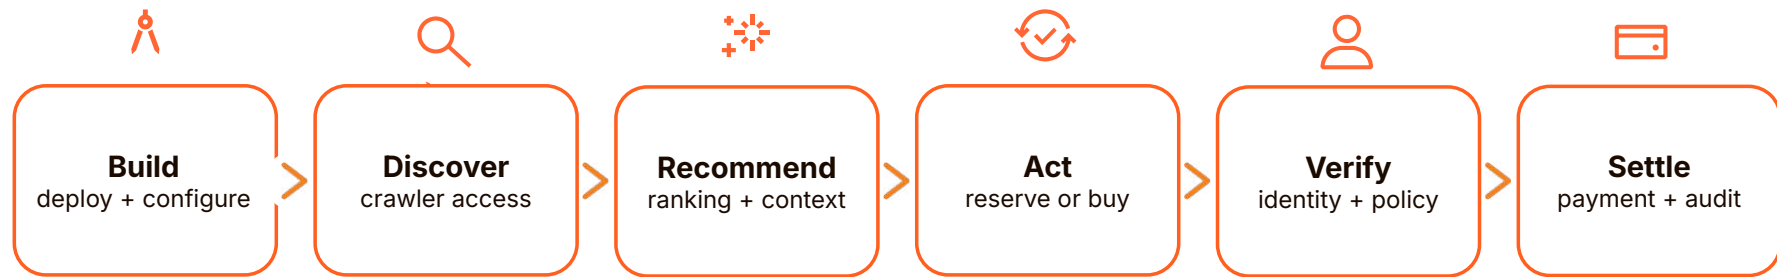
Any access request **can be combined with a payment request.**

We have the full stack for agentic commerce



Every agent request can become a metered event

The economic surface expands from content access to recommendations, payments, fulfillment, and audit.



The strategic shift: from protecting websites from bots to monetizing trusted automated demand



It's already
happening

Large media companies use Cloudflare to do deals



People Inc.

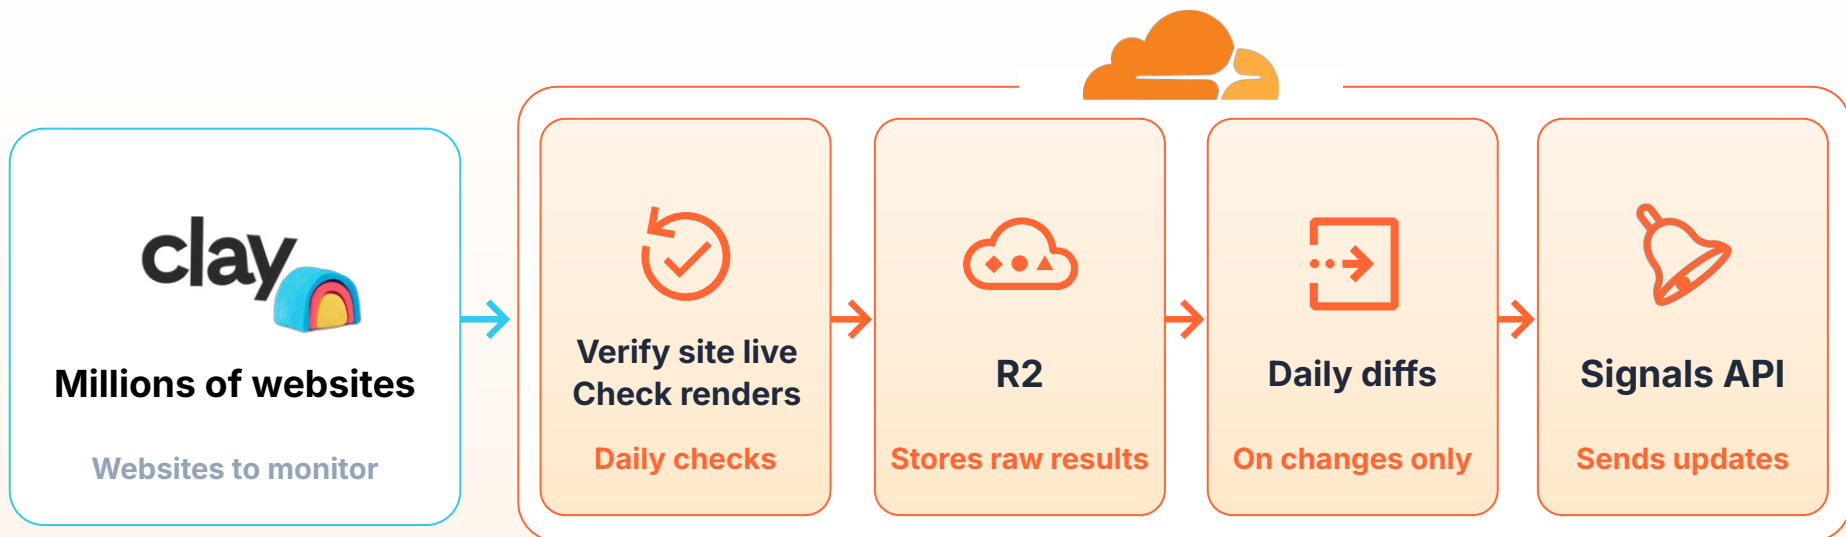
Ziff
Davis

A
The Atlantic

People Inc. uses the technology from web infrastructure provider Cloudflare to block other AI crawlers, prompting AI players to approach it with content deals.

In September 2025, People attributed its decision to leverage Cloudflare's technology as a way to get AI companies to the negotiating table, noting that its progress on deals was "much further along" after adopting the solution.

Startup AI companies use Cloudflare to access the Internet



Proprietary crawl intelligence lets us price niche content programmatically

Crawl Attempts	Sample Publisher	Sample Health Site
AI Content Scraper	23 MM	112 K
AI Chat User	120 MM	787 MM
Search Bot	24 MM	459 K
AI Agent	2 MM	12 K
Search Bot	211 MM	70 MM



Niche Content Scraping

AI crawlers are scraping specific content (e.g., health and science and local news)

We are just getting started

Media sector **today**

37%+ revenue growth

in Q1'26 vs. Q1'25

60%+ expansion rate

for retained customers in Q1'26

1.3x higher revenue per customer

than company average

~117% DNR

in Q1'26

Tomorrow...



Content Marketplace

Programmatic monetization of content at scale



Signals

Standardized, machine-readable vocabulary for real-time decision-making



Payments

Micropayment rails for crawlers, autonomous agents, and agentic commerce at scale



"Easy Button" for Agents

Smart infrastructure with built-in guardrails

Cloudflare was built for this moment

Act I + Act III = Act IV



Connect & Protect



Build



Agentic Internet

Go-To-Market



Mark Anderson

President of Revenue


What customers are saying

The market has come to Cloudflare



The market **has come to Cloudflare**

Then (2024)




Brand Awareness Trend

Phase: Market Education

- "Who is Cloudflare?"
- "Why Cloudflare?"
- "Why now?"



Now (2026)



\$1M+ Large Customers

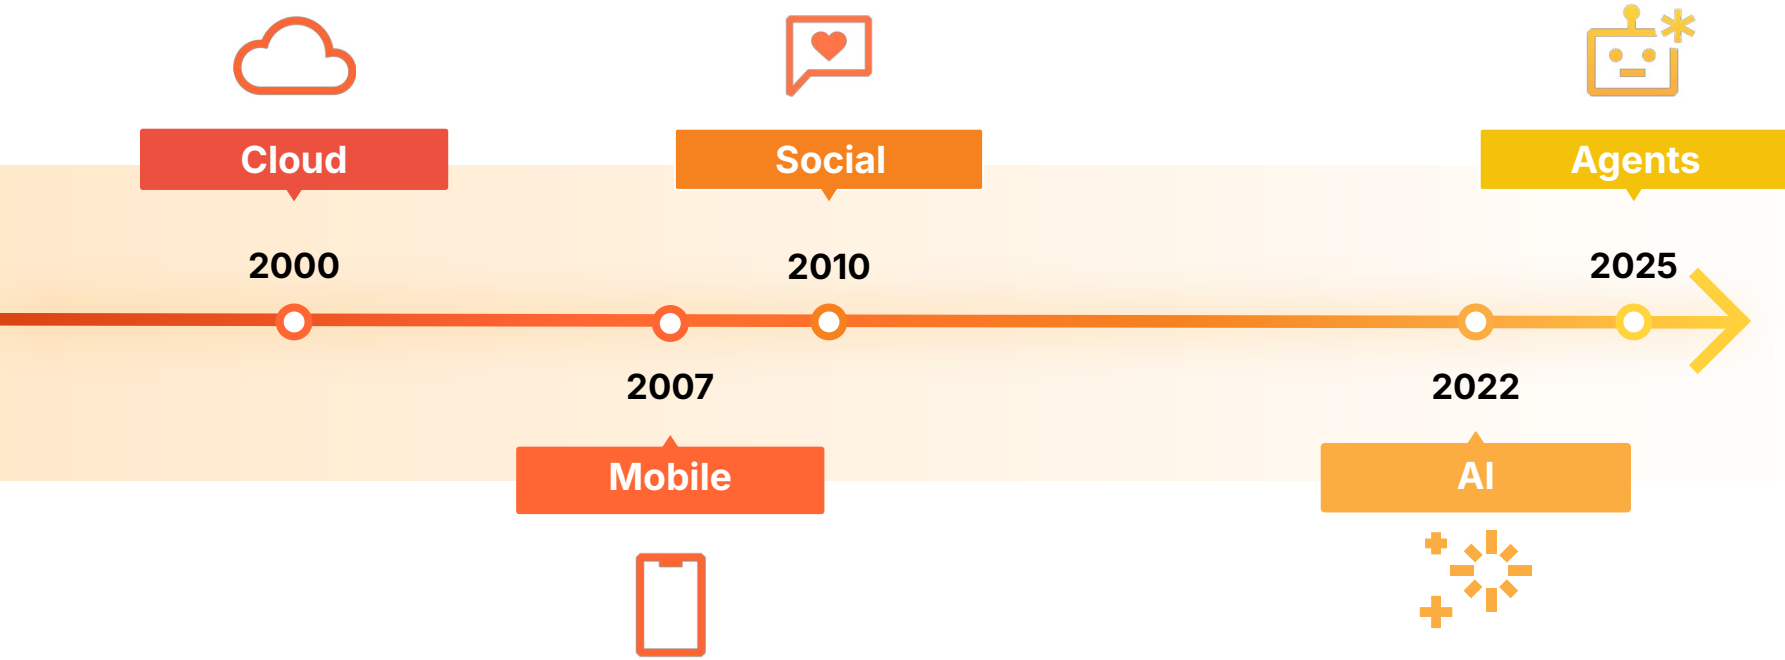
Phase: Market Pull

- AI modernization
- Security transformation
- Infrastructure consolidation

Source: Cloudflare Brand Awareness Study. \$1M customers are Cloudflare Customers. See Appendix for how we define "Large Customers," which are the same as \$100K+ customers. \$1M+ customers are defined in the same manner, except we use \$1M+, respectively, of Annualized Revenue instead of \$100K+.

Enterprise adoption continues to accelerate

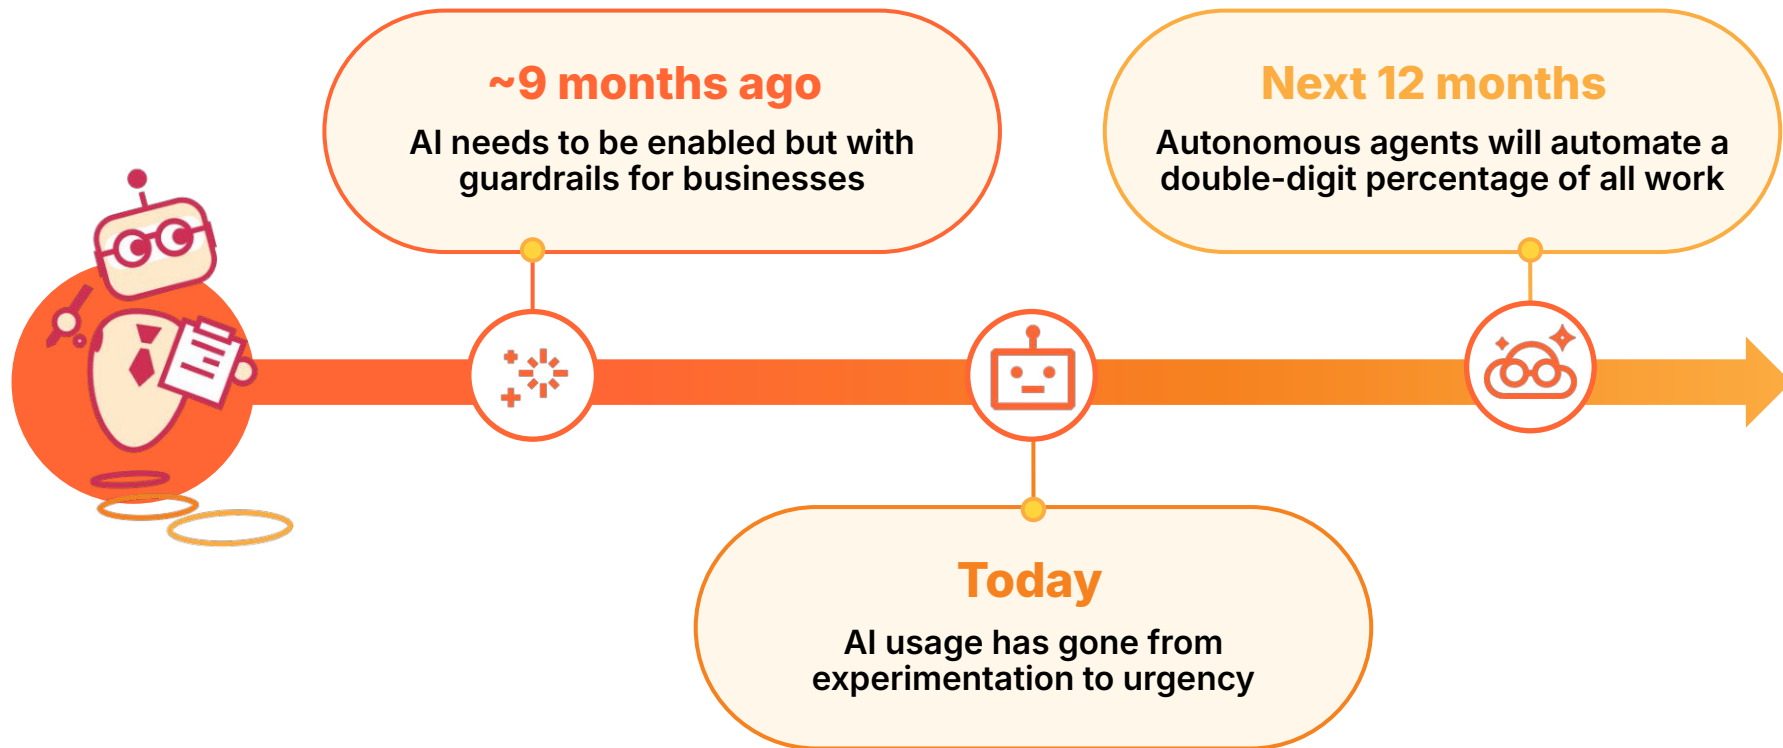
We are in the middle of a platform shift



AI and
agents
change
everything




AI has moved from experimentation to **urgency**



The agentic surge is real


1,700%+
INCREASE

in daily AI agent requests
on Cloudflare's network
from June 1, 2025
to May 31, 2026



Frontier AI models change the attack landscape

Average time between vulnerability disclosure → confirmed exploitation is **collapsing**




Project Glasswing: what Mythos showed us

2024-05-18

Grant Bourzkas

9 min read

This post is also available in [简体中文](#), [Français](#), [Deutsch](#), [Italiano](#), [日本語](#), [한국어](#), [Português](#), [Español](#), [Nederlands](#) and [繁體中文](#).



For the last few months, we've been testing a range of security-focused LLMs on our own infrastructure. These LLMs help identify potential vulnerabilities in our own systems, so we can fix them – and they also show us what attackers are going to be able to do with the latest models.

**Customers
need a
modern
platform for
an AI-driven
world**

Everyone is
building &
running
agents

Frontier AI
models are
changing
the attack
landscape

IT budgets and headcounts are only **marginally increasing**...

2.8%


average **IT budget**
growth in 2026

1.3%

average **tech worker**
headcount growth
in 2026

...but CIOs are **expanding AI spending significantly**

Percent of organizations increasing investment in 2026 digital priorities




IT leaders' focus initiatives: accelerate AI adoption with security built-in



Secure workforce use of AI

Discover and control how employees interact with AI tools (e.g., ChatGPT, Claude)



Govern AI agents

Secure interactions between human users and AI agents



Protect applications

Defend applications and websites — including those with AI endpoints (e.g., chatbots) — against data loss and attacks



Build AI securely

Build AI applications and AI agents for internal or customer use

Customers don't need a whole new paradigm to build, adopt, & secure AI.

They need a platform that was built for this moment in the first place.



for users

























& for agents



**Cloudflare was
built for this
moment.**

Cloudflare is the **only vendor** to help customers **connect, protect, and build AI**


We offer **AI security and AI infrastructure products** — all on one platform and one network.

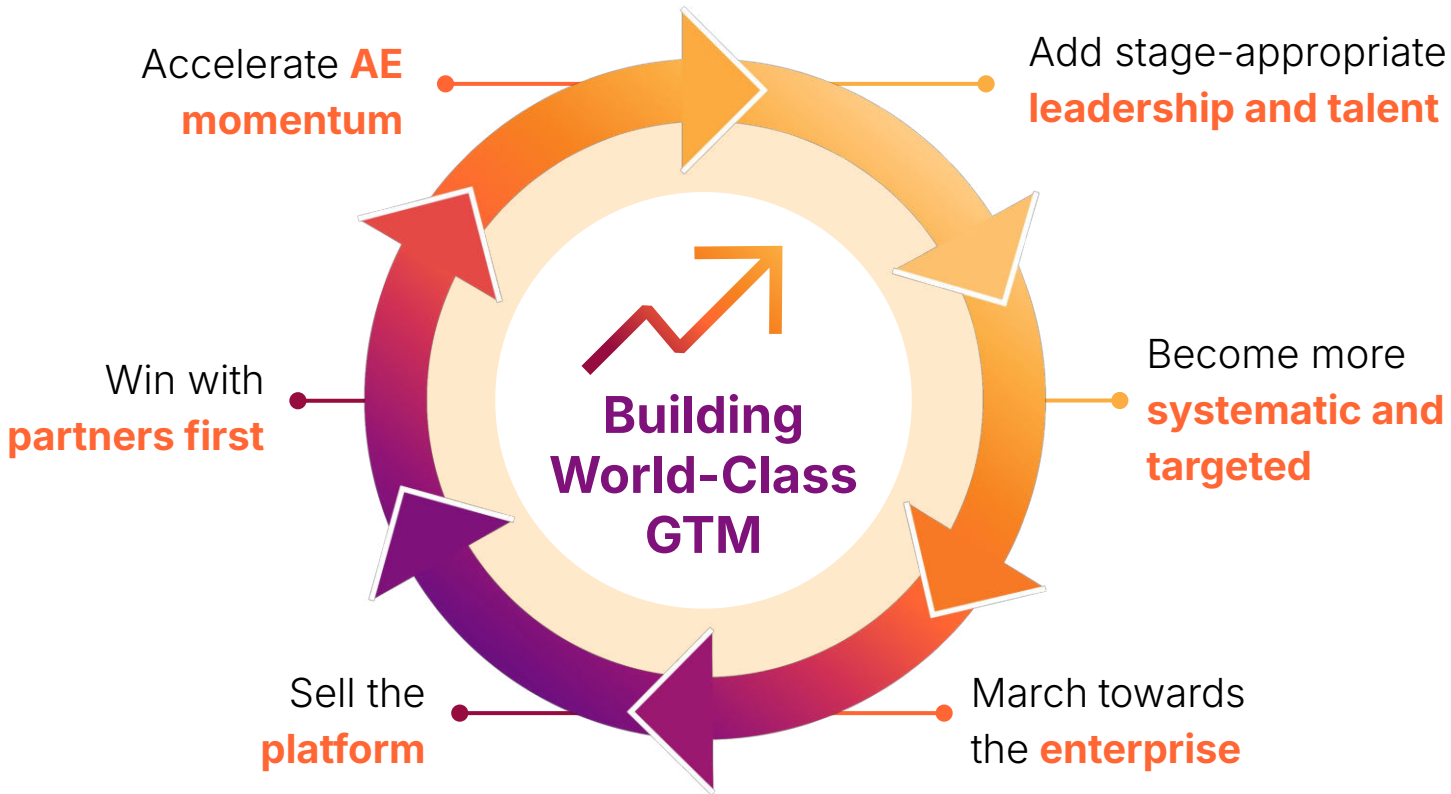
	 CLOUDFLARE			
Category	Secure workforce use of AI	Govern AI agents	Protect AI-powered applications	Build AI securely
Cloudflare products	Cloudflare One	Remote MCP servers, Access with MCP Server Portals	Application Security (inc. AI Security for Apps)	Workers AI, AI Search, AI Gateway, Agents SDK, Dynamic Workers, Containers, Artifacts
Other major vendors	SASE / SSE      	Developer-focused   Hyperscalers  	App Security / WAAP  	Hyperscalers   
Specialized startups	  	 	 	
Cloudflare advantages	Comprehensive AI lifecycle protection			Future-proof global architecture
			Model-agnostic deployment	

 For CIOs & CISOs

 For CTOs & Developers

Building a **world-class** go-to-market organization





Established new GTM leadership team *and* upleveled sales management and the field



Jeff Samuels
Chief Marketing Officer



Rick Congdon
GEO VP, Americas



Tony Van den Berge
GEO VP, EMEA



Goran Ristic
GEO VP, APAC



Sayoko Matsumoto
GEO VP, Japan



Aly Cabral
VP, Developer GTM



Mike Ogden
VP, Global Revenue Strategy & Operations



Jeff Ward
VP, Global Sales and Renewals



Tom Evans
Chief Partner Officer



Gregory Marchwinski
VP, Global Customer Services


61%

of GTM
managers
are new to
Cloudflare
since Q1 2024

68%

of account
executives
are new to
Cloudflare
since Q1 2024


Improved go-to-market approach to best address the unique needs of our diverse customer set



Defining clear market segments allows us to apply specific sales models to **target customers' challenges and opportunities** within each segment.

Enterprise momentum continues to build...


\$1M+ Customers vs. Annualized Revenue



\$1M+ customers represent **less than 1%** of Cloudflare's paying customers.

Source: Latest company reports, filings, and investor presentations. Annualized Revenue = quarterly revenue x 4, except for UiPath and Splunk displaying Annual Recurring Revenue. Chart uses data from CY-2020 to CY-2025, except for ServiceNow (2015 to 2018), Salesforce, Splunk, and Palo Alto Networks whose current \$1M customer counts exceed the bounds of the chart. Salesforce, CrowdStrike, DataDog and MongoDB quarterly customer count extrapolated linearly from Q4 annual customer disclosures.

...across all industries



Connect, protect, and build *everywhere*

Connect and protect
external-facing systems

 Cloudflare Network Services

 Cloudflare Application Services

Connect and protect
internal-facing systems


 Cloudflare Zero Trust Services

Build the next-generation of applications


 Cloudflare Developer Platform

One Platform  One Network

The power of **partnership**



Partners as a % of revenue



Physics 101

$$p = m v$$

Momentum, n.

the quantity of motion in a moving body, expressed as the product of its mass and velocity

Mass, n.


the quantity of matter within a given object

Velocity, n.

a quantity that designates speed together with the direction of movement

Productivity


Rolling 4-Quarter ACV per Account Executive (Net)




We are encouraged by the continued improvement in **sales productivity** to new all-time highs.

Increased hiring of stage-appropriate GTM talent

Account Executive Capacity (Net) Y/Y Growth



Composition of New AE Hires



We continue to increase hiring in sales — with a focus on **enterprise account executives** with proven track records.

What's next for 2026 and beyond?



We are our
own **most**
demanding
AI customer




AI agents are becoming **full-cycle co-sellers**

Across the industry, sales reps are being stretched thin by two opposing forces: **growing customer demands and limited time to deliver.**

AI agents are increasingly taking ownership of time-intensive tasks, allowing AEs to focus more time on higher-value, customer-facing work:

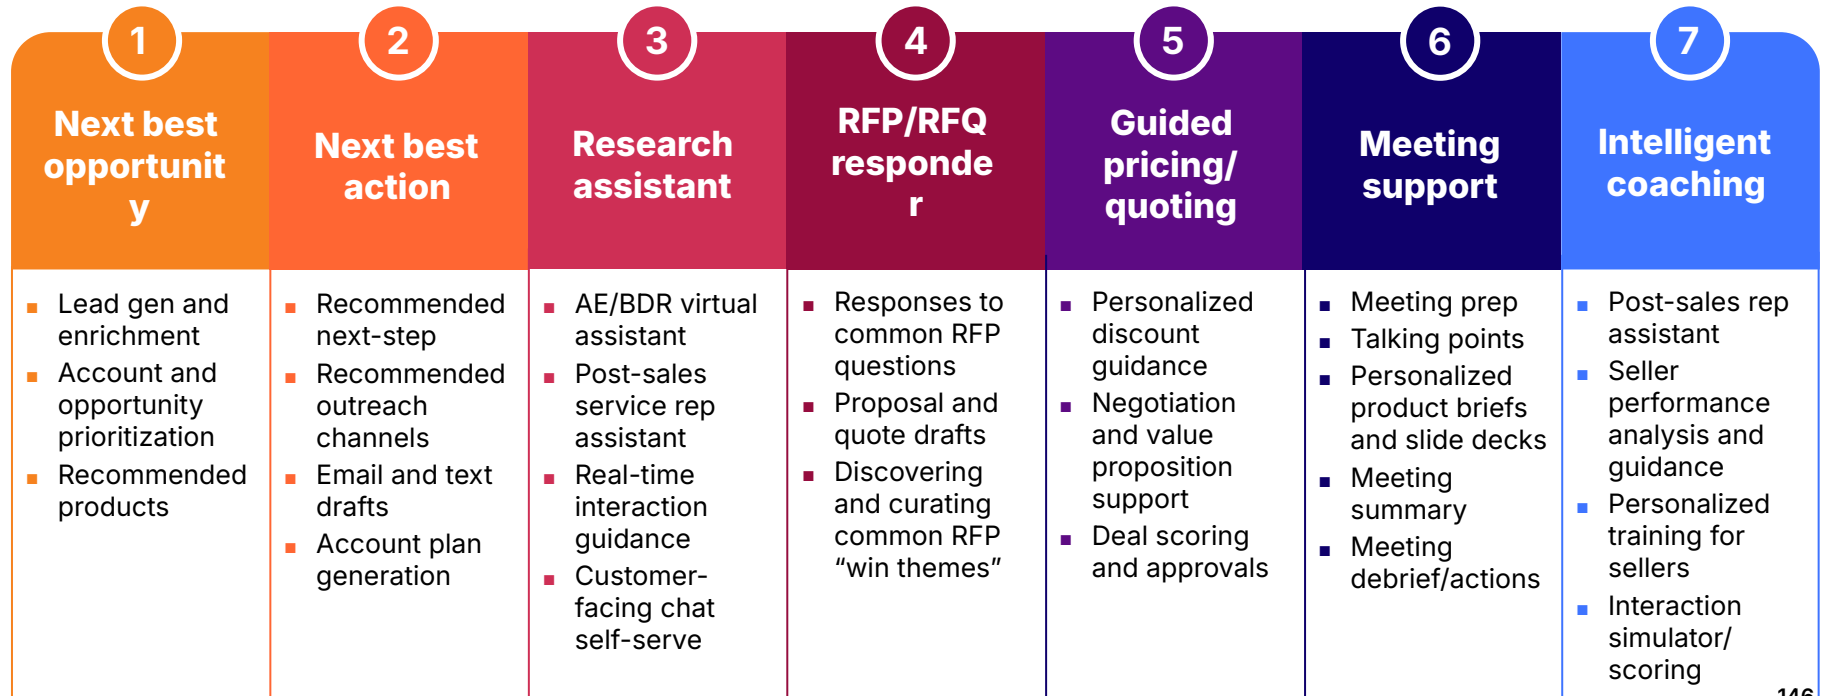
- ✓ **growing pipeline**
- ✓ **closing deals**
- ✓ **increasing revenue**

Salesforce: The State of Sales Report – “How Reps Spend Their Time During an Average Workweek”



AI-powered tools are **transforming** sales strategies

Examples of AI use cases from one end of the sales process to the other:



The customer lifecycle reimaged with AI



1. Onboarding

Setting the foundation

Standardizing the initial journey for rapid time-to-value and technical integration.

- Kick Off
- Success Plan Builder



2. Adoption

Proving value

Driving deep platform usage through strategic business reviews and automation.

- Consumption Monitoring
- QBR Automation



3. Renewal

Solidifying the relationship

Ensuring commitment and identifying strategic expansion opportunities.

- Expansion Signals
- DNR Tracking



Automation & Tooling

- Overview of Purchase Deck
- Pre Call Prep Agent
- Post Call Follow-up Agent
- EBR Creator & Auto-Scheduling
- Overage & Success Triggers
- Incoming Request Agent
- Quoting & Expansion Tool
- 90-Day Renewal Notification
- Growth Projection Proposal

Go-to-market
at Cloudflare
has achieved
**escape
velocity**



Finance



Thomas Seifert


Chief Financial Officer

Traffic growth is accelerating


1,700% + INCREASE

in daily AI agent requests
on Cloudflare's network
from June 1, 2025
to May 31, 2026

Millions of Requests per Second on Cloudflare's Network (increase in Q1'25 vs. Q1'26)



Cloudflare Revenue Models 101



Understanding Cloudflare's revenue models


1 Flat Rate (Free, Pro, and Business plans)	2 Pure Usage (Pay-As-You-Go)	3 Pool of Funds (Enterprise committed value contracts)	4 Cap Model (Enterprise committed capacity contracts)
<p>One fixed price for unlimited/uncapped data transfer, restricted by feature access.</p> <p>CUSTOMER PERSPECTIVE</p> <p>Pros: Absolute budget certainty.</p> <p>Cons: Rigid feature limits; paying for unused headroom if site traffic is low.</p> <p>CLOUDFLARE PERSPECTIVE</p> <p>Pros: Highly predictable, automated self-serve recurring revenue; low sales overhead.</p> <p>Cons: Constrained revenue upside as customer usage increases.</p>	<p>Billed in arrears for the exact number of units or resources consumed.</p> <p>CUSTOMER PERSPECTIVE</p> <p>Pros: No upfront financial commitment required to build.</p> <p>Cons: Budget volatility.</p> <p>CLOUDFLARE PERSPECTIVE</p> <p>Pros: Highly competitive developer entry point.</p> <p>Cons: More unpredictable revenue forecasting.</p>	<p>Customer commits a budget value upfront to draw down flexibly across Cloudflare services consumed.</p> <p>CUSTOMER PERSPECTIVE</p> <p>Pros: Architecture flexibility; allows deployment of new products instantly without procurement friction.</p> <p>Cons: Requires upfront monetary commitment; risk of losing funds if underutilized by contract end.</p> <p>CLOUDFLARE PERSPECTIVE</p> <p>Pros: Significantly increases Customer Lifetime Value (LTV); secures massive, multi-year upfront contract value.</p> <p>Cons: Requires account management and sales engineering to ensure product utilization.</p>	<p>Fixed base recurring fee including a monthly entitlement, with tiered overages beyond the cap.</p> <p>CUSTOMER PERSPECTIVE</p> <p>Pros: Combines a safe, predictable budget floor with room to scale.</p> <p>Cons: Incurs overage expenses if standard legitimate business growth rapidly outpaces the contracted cap.</p> <p>CLOUDFLARE PERSPECTIVE</p> <p>Pros: Secures a high, predictable ARR floor while establishing a friction-free, high-margin revenue expansion engine.</p> <p>Cons: Requires continuous monitoring and "cap increase" negotiations for accounts consistently exceeding caps.</p>

Monthly cap contracts provide **predictable** revenue for Cloudflare and cost for customers...

Illustrative Example of Cap Model

Why customers select a cap model:


- ✓ offers predictable budgeting, as long as the company operates within its growth projections (which is the basis of the caps in the contract)
- ✓ provides flexibility to raise or lower caps as business conditions change



The majority of Cloudflare's contracts are based on a cap model with the requests being the most significant unit of measure for revenue.

...but accelerating traffic growth can result in **earlier-than-expected increases to cap model contracts**


Illustrative Example of Cap Model with Accelerated Usage



When actual usage outpaces a customer's initial forecasts, the cap model triggers early "upsizing" of the committed capacity — **moving revenue growth earlier** than planned *but* with a lag to the inflection in usage.

A Leading Technology Company: Increased monthly cap twice in 2025


Example:
Existing customer increasing monthly cap contract



Note: FY2023 and FY2024 revenue reflects actual full fiscal-year revenue; Q1'25-Q1'26 revenue is annualized (quarterly revenue x 4). Overage is calculated by applying the customer's contracted data-transfer overage rates to billable data transfer above the monthly cap, annualized on the same basis (FY2024 shown as actual). Billable data transfer is the peak monthly billable data transfer (TB) within each period, plotted against the monthly data-transfer cap.

Pool of funds accelerate platform adoption...

Pool of Funds as a % of New ACV Booked




Enterprises are **increasingly accustomed** to purchasing cloud services through flexible consumption models.


...but can impact the "shape" of financial reporting

Revenue recognition is often **lower in the near term** but **accelerates over course of the contract** as customers accelerate adoption and expansion.

2024 PoF Cohort: PoF Revenue vs. If-Ratable Revenue




2025 PoF Cohort: PoF Revenue vs. If-Ratable Revenue



Pool of funds contracts are effectively ramped deals.

A Leading AI Company: Signed two PoF deals in 2024 and 2025


Example:
Existing customer transitioning to pool of funds



Note: 2022-2024 reflect full-year actual revenue; Q1'25-Q2'26 reflect annualized run-rate revenue shown on a common annual-dollar basis. Q2'26 run-rate is annualized from April and May 2026 actuals, weighted toward the most recent month. Y/Y and Q/Q growth are calculated on these annualized figures.

Drivers of Value Creation

Growth + Profitability





Growth


Drivers of high, sustainable growth




Expanding TAM



Growing pipeline



Improving sales productivity and growing capacity



Building momentum with larger customers




Increasing platform adoption



Accelerating traction in key product areas

A consistently expanding TAM




Areas for Potential Incremental Growth

- Act IV
- Database
- Internet of Things
- 5G Cellular
- Network Services

Note: Calculations performed by Cloudflare. Graphic created by Cloudflare based on Gartner® research. Figures reflect the markets addressed by our current product offerings. For full citation details on the Gartner forecasts used in Cloudflare's estimations, see the Appendix.

Pipeline continues to grow


Pipeline (\$M)




Robust pipeline growth gives us confidence to invest in sales capacity.

Go-to-market initiatives are paying off

Rolling 4-Quarter ACV per Account Executive (Net)



Account Executive Capacity (Net) Y/Y Growth




Actual Forecast

We continue to invest in quota-carrying sales **capacity** and expect that **productivity** to continue to grow for a very long time.

Large customers are getting larger

Revenue (\$M) by All Customers




42%
of the Fortune 500


Note: Growth factors (e.g., 6x for <\$100K) are calculated as Q4'25 revenue divided by Q1'18 revenue. Fortune 500 figure is as of March 31, 2026.

Land-and-expand fuels platform adoption

Annualized Recurring Revenue by Product SKU



Annualized Recurring Revenue by Product Attach Rate Cohorts (\$M)



Our **success in driving platform adoption** is evidenced by the consistently increasing amount spent with Cloudflare in higher product attach rate cohorts.

Note: Annualized Recurring Revenue by Product SKU as of December 31, 2025 for contracted customers. Annualized Recurring Revenue by product attach rate cohorts is as of the fourth quarter of each respective year. See the Appendix for how we define Annualized Recurring Revenue (ARR) and contracted customers.

Cloudflare One

↑ 43%

y/y growth in ARR in
2025

Developer

↑ 137%

y/y growth in ARR in
2025

Cloudflare's
multi-act play
is only
**just getting
started.**



Profitability

Unit economics

Cost to Book (CTB)

What does it “cost to book” \$1 of incremental ARR? CTB = Sales & Marketing cost per dollar of incremental ARR

Cost to Serve (CTS)

How much does it “cost to serve” the lifetime revenue? CTS = Delivery, R&D, Support & G&A cost as % of ARR


Attrition

What’s the expected lifetime revenue from \$1 ARR? Lifetime Revenue (LTR) = \$1 ARR / attrition rate

Unit Economic Margin

What is the economic margin created by \$1 ARR? Unit Economic Margin = $[LTR - CTB - (CTS * LTR)] / LTR$


Unit Economic Margin by Component



Gross margin reporting is influenced by multiple factors

- The percentage of paid versus free traffic on our network, which is neutral to operating margin, has grown significantly.
- Workers developer platform products, which are currently not optimized for gross margin (although gross margin is improving), continue to deliver outsized growth.

Change in Gross Margin by Factor (since Q1'24)



While gross margin may continue to trend down in the near-term from these dynamics, the scalability, elasticity, and efficiency of our network remain unchanged.

How revenue mix influences unit economic KPIs




Act I



Act II




Act III



Act IV

UNIT ECONOMICS DEFINE THE GROWTH/MARGIN FRONTIER



$$\text{VALUE} = \frac{\text{REVENUE} - (\text{CTS}) - \text{CTB} \times (\text{ATTRITION} + \text{GROWTH})}{\text{REVENUE}}$$


UNIT ECONOMIC MARGIN = 44%




UNIT ECONOMICS BY ACT

ACT I

- COST TO BOOK : AVERAGE
- COST TO SERVE : AVERAGE
- ATTRITION : HIGHER



UNIT ECONOMICS BY ACT



ACT II

- COST TO BOOK : HIGHER
- COST TO SERVE : LOW
- ATTRITION : LOWER

Lower gross margin **doesn't** mean lower unit economic margin — especially for platform businesses at scale

CASE STUDY FOR ACT III

Microsoft Intelligent Cloud (\$B)


	FY2015	FY2020	FY2025
Revenue y/y growth	9.1%	24.1%	21.5%
-- Azure % of IC revenue	~4.0%	~29.0%	~71.0%
Gross margin	76.1%	69.1%	62.2%
OPEX % of IC revenue	34.5%	31.3%	20.2%
Operating margin	41.6%	37.9%	42.0%
Rule of 40	50.7%	61.9%	63.5%

COST TO BOOK: Early customer acquisition relies on efficient product-led growth (PLG), self-serve onboarding, and developer advocacy. Customer acquisition then expands to institutional enterprise cross-selling, translating to high sales productivity and a lower blended cost to book per enterprise dollar.

COST TO SERVE: Although initially higher early in the adoption lifecycle, extreme operating expense leverage in R&D and G&A is achieved as platform revenue scales, lowering aggregate cost to serve (even at lower gross margin).


ATTRITION: Operational friction and expensive platform switching costs create high retention rates.

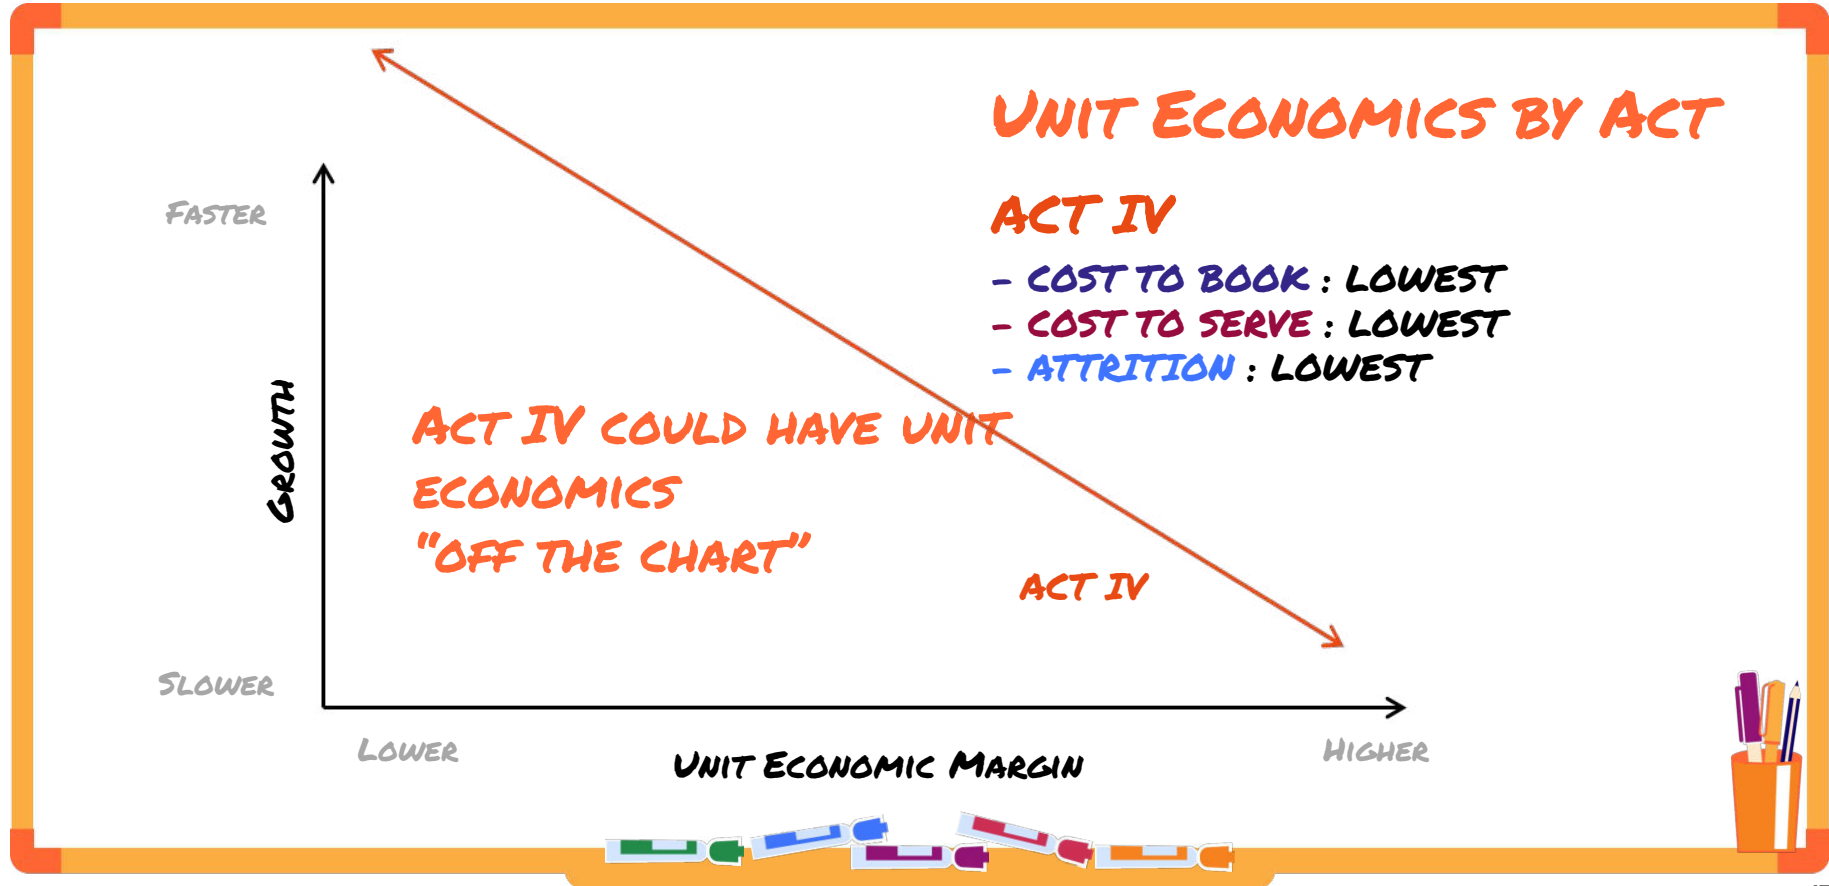
At scale, the compounding operating leverage of the platform business model drives exceptionally strong unit economic margins.



PLATFORM BUSINESSES
POSSESS THE UNIQUE
CAPACITY TO
STRUCTURALLY SHIFT
THEIR GROWTH/MARGIN
FRONTIERS UPWARD AND
TO THE RIGHT AS
REVENUE SCALES.







What are our financial targets?



Before year-end 2028

\$5B+

annualized revenue

Year-end 2025

\$2.4B

annualized revenue

Our "Rule of" North Star...

Since IPO

40%

2027

50%

2028+

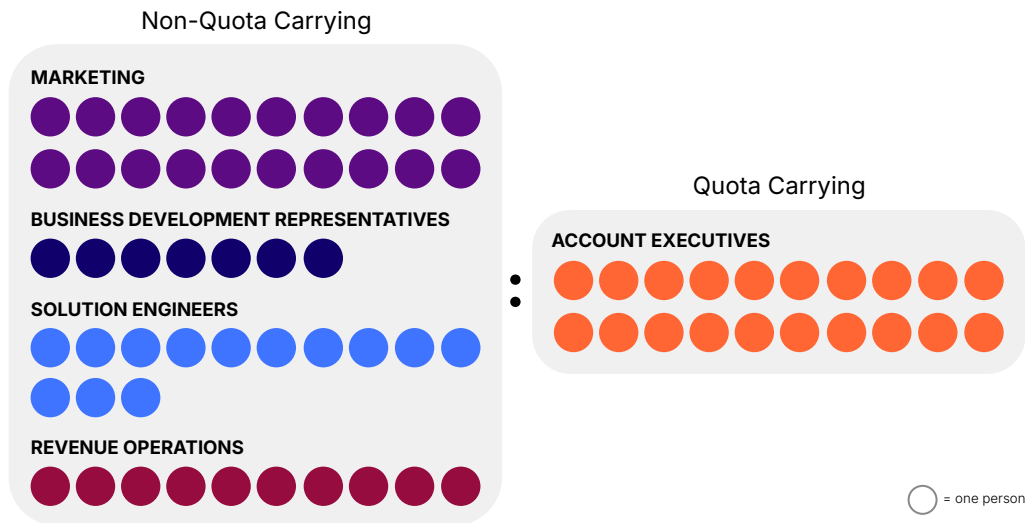
50%+

AI is driving **better** cost to book...

Every GTM metric across the technology industry has historically scaled based on a ratio of employees who carry quota (AEs) to every other support function within the GTM organization (marketing, BDRs, SEs, RevOps, etc.).

AI is fundamentally changing these long-standing ratios, enabling a **higher ratio of quota-carriers** within the GTM organization.

Illustrative Example of Ratio-Oriented GTM Model



AI augments traditional GTM playbooks — automating more support functions, driving higher conversion across the funnel, and allowing AEs to own more of the workflow.

...that can be reinvested in **more AEs** to produce **higher growth**

Illustrative Example of Ratio-Oriented GTM Model — Status Quo vs. with AI Automation

	----- Status Quo -----			----- w/ AI Automation -----		
	Headcount	AE:[x] ratio	Total Cost	Headcount	AE:[x] ratio	Total Cost
AEs	30		\$7,500,000	40		\$10,124,654
BDRs	10	3:1	\$850,000	10	4:1	\$860,596
SEs	20	3:2	\$4,300,000	13	3:1	\$2,902,401
RevOps	15	2:1	\$2,625,000	13	3:1	\$2,362,419
Marketing	30	1:1	\$3,000,000	20	2:1	\$2,024,931
Total	105		\$18,275,000	98		\$18,275,000
	S&M Cost =		\$18,275,000	S&M Cost =		\$18,275,000
	ACV =		\$26,250,000	ACV =		\$35,436,288

↑ 35%

From one end of the sales process to the other, AI agents are augmenting and automating support functions within GTM — **freeing budget for more quota-carrying AEs.**

Supercharging a 10x engineer into a 1000x engineer with AI

ONE ENGINEER • ONE WEEK

Reimplementing Next.js

Steve Faulkner, Sr. Director of Cloudflare Workers (managing 80+ engineers), reimplemented Next.js from scratch using AI coding agents with max thinking enabled.

The result was vinext, a drop-in replacement for Next.js, built on Vite, that deploys to Workers with a single command. In early benchmarks, vinext builds production apps up to 4x faster and produces client bundles up to 57% smaller.

The whole process cost about \$1,100 in tokens and took less than one week.

TWO ENGINEERS • TWO MONTHS

Rebuilding WordPress

One product manager and one engineer at Cloudflare used AI agents over two months to rebuild WordPress from scratch in TypeScript.

The result was EmDash, a serverless, AI-native CMS built on Astro and Cloudflare Workers. Plugins run in sandboxed isolates, solving the security flaw behind 96% of WordPress vulnerabilities, and every site has built-in x402 payments for the agentic web.

A rebuild of a 23-year-old platform, which is the content management system for 40%+ of the Internet, took two months.

EVERY MERGE REQUEST • EVERY REPO

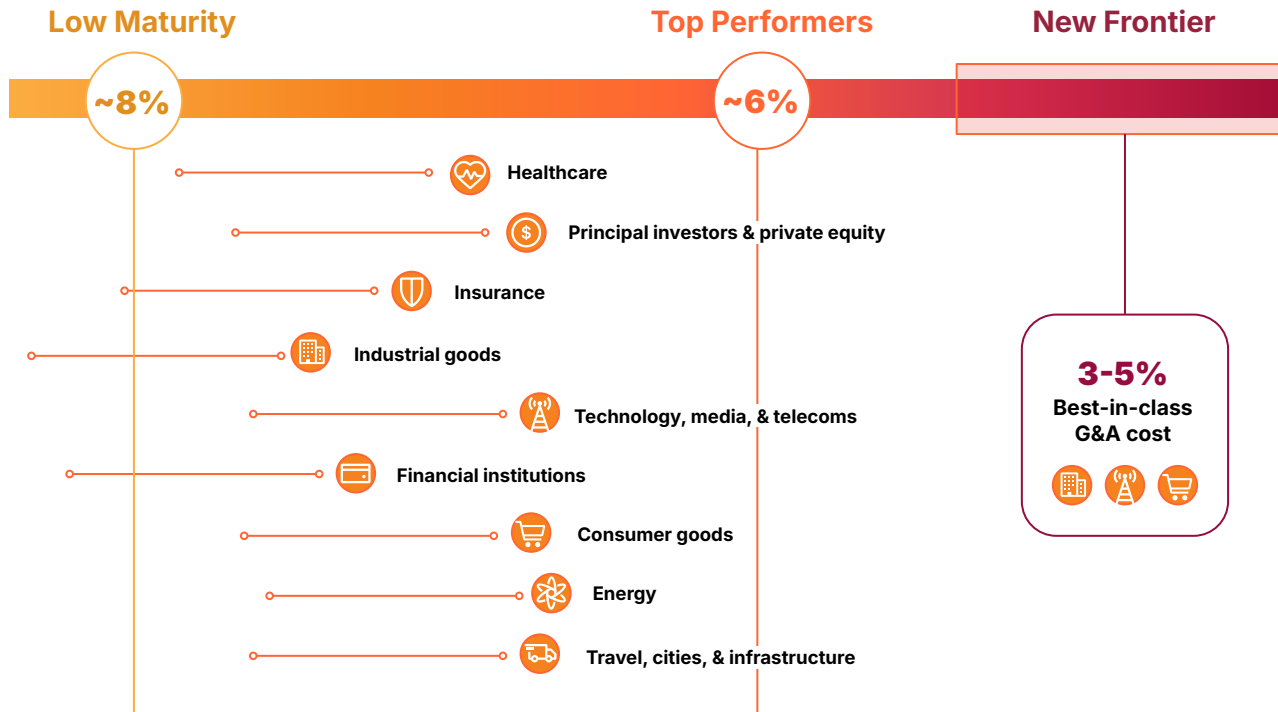
Reviewing Code across Every Repo

Every merge request on Cloudflare's standard CI pipeline gets an AI code review, running up to seven specialized agents — security, performance, code quality, documentation, compliance — coordinated by a top-tier model that posts a single structured review.

In the first 30 days, the system ran 131,246 reviews across 48,095 merge requests in 5,169 repos, with a median completion time of 3 minutes 39 seconds.

Median cost: \$0.98 per review. 100% merge request coverage. Built on Workers and AI Gateway.

AI is transforming G&A workflows



G&A is being rebuilt around AI-augmented workflows

- **Finance & Accounting:** continuous close, automated reconciliations, real-time variance commentary
- **HR & Recruiting:** candidate screening, internal policy Q&A, performance-review synthesis
- **Legal & Compliance:** contract review and redlines, regulatory change monitoring (days to minutes)
- **IT & Knowledge Management:** service-desk ticketing and deep research are scaled agentic use cases
- **Procurement:** sourcing intake, RFP drafting, contract-clause comparison

Long-term operating model

	2022	2023	2024	2025	Prior Long-Term Model	Updated Long-Term Model
Gross Margin	78 %	78 %	79 %	76 %	75% - 77%	70% - 77%
Sales & Marketing (% of revenue)	42 %	40 %	38 %	36 %	27% - 29%	23% - 25%
Research & Development (% of revenue)	19 %	17 %	16 %	16 %	18% - 20%	15% - 17%
General & Administrative (% of revenue)	14 %	12 %	11 %	10 %	8% - 10%	3% - 5%
Operating Margin	4 %	9 %	14 %	14 %	20%+	30%+
Free Cash Flow Margin	(4) %	9 %	10 %	12 %	~25%	30 - 35%+

Achieve GAAP profitability by 2028 latest

Q&A



Matthew Prince

Co-founder & Chief
Executive Officer



Michelle Zatlyn

Co-founder &
President



Thomas Seifert

Chief Financial Officer



Mark Anderson

President of
Revenue

Appendix

GAAP to Non-GAAP Reconciliation (Annual)

	2021	2022	2023	2024	2025
GAAP cost of revenue	\$147,134	\$232,610	\$307,005	\$378,702	\$552,525
Less: Stock-based compensation expense & related payroll taxes	-\$3,703	-\$6,770	-\$8,360	-\$11,597	-\$14,605
Less: Amortization of acquired intangible assets	-\$2,946	-\$13,444	-\$17,702	-\$11,084	-\$13,172
Non-GAAP cost of revenue	\$140,485	\$212,396	\$280,943	\$356,021	\$524,748
GAAP gross profit	\$509,292	\$742,631	\$989,740	\$1,290,924	\$1,615,412
Add: Stock-based compensation expense & related payroll taxes	\$3,703	\$6,770	\$8,360	\$11,597	\$14,605
Add: Amortization of acquired intangible assets	\$2,946	\$13,444	\$17,702	\$11,084	\$13,172
Non-GAAP gross profit	\$515,941	\$762,845	\$1,015,802	\$1,313,605	\$1,643,189
GAAP gross margin	78%	76%	76%	77%	75%
Non-GAAP gross margin	79%	78%	78%	79%	76%
GAAP sales and marketing expense	\$328,065	\$465,762	\$599,117	\$745,791	\$920,817
Less: Stock-based compensation expense & related payroll taxes	-\$32,869	-\$53,692	-\$76,711	-\$95,763	-\$137,848
Less: Amortization of acquired intangible assets	-	-\$1,725	-\$2,300	-\$1,663	-\$1,826
Less: Acquisition-related and other expenses	-	-\$265	-	-	-
Less: One-time compensation charge	-	-	-	-\$15,000	-
Non-GAAP sales and marketing expense	\$295,196	\$410,080	\$520,106	\$633,365	\$781,143
GAAP S&M expense as a % of revenue	50%	48%	46%	45%	42%
Non-GAAP S&M expense as a % of revenue	45%	42%	40%	38%	36%
GAAP research and development expense	\$189,408	\$298,303	\$358,143	\$421,374	\$512,489
Less: Stock-based compensation expense & related payroll taxes	-\$61,056	-\$112,277	-\$140,074	-\$151,936	-\$174,622
Less: Acquisition-related and other expenses	-	-\$3,682	-	-	-
Non-GAAP research and development expense	\$128,352	\$182,344	\$218,069	\$269,438	\$337,867
GAAP R&D expense as a % of revenue	29%	31%	28%	25%	24%
Non-GAAP R&D expense as a % of revenue	20%	19%	17%	16%	16%
GAAP general and administrative expense	\$119,503	\$179,769	\$217,965	\$278,520	\$389,311
Less: Stock-based compensation expense & related payroll taxes	-\$19,706	-\$45,027	-\$62,355	-\$97,127	-\$162,863
Less: Acquisition-related and other expenses	-\$380	-	-	-\$702	-\$3,909
Less: Lease impairment charges	-	-	-	-	-\$5,097
Less: Legal reserve and settlements	-	-	-	-	\$2,886
Non-GAAP general and administrative expense	\$99,417	\$134,742	\$155,610	\$180,691	\$220,328
GAAP G&A expense as a % of revenue	18%	18%	17%	17%	18%
Non-GAAP G&A expense as a % of revenue	15%	14%	12%	11%	10%
GAAP loss from operations	-\$127,684	-\$201,203	-\$185,485	-\$154,761	-\$207,205
Add: Stock-based compensation expense & related payroll taxes	\$117,334	\$217,766	\$287,500	\$356,423	\$489,938
Add: Amortization of acquired intangible assets	\$2,946	\$15,169	\$20,002	\$12,747	\$14,998
Add: Acquisition-related and other expenses	\$380	\$3,947	-	\$702	\$3,909
Add: One-time compensation charge	-	-	-	\$15,000	-
Add: Lease asset impairment expense	-	-	-	-	\$5,097
Add: Legal reserve and settlements	-	-	-	-	-\$2,886
Non-GAAP loss from operations	-\$7,024	\$35,679	\$122,017	\$230,111	\$303,851
GAAP operating margin	-19%	-21%	-14%	-9%	-10%
Non-GAAP operating margin	-1%	4%	9%	14%	14%

GAAP to Non-GAAP Reconciliation (Quarterly)

	Q1 2025	Q2 2025	Q3 2025	Q4 2025	Q1 2026
GAAP cost of revenue	\$115,576	\$128,677	\$146,316	\$161,956	\$184,158
Less: Stock-based compensation expense & related payroll taxes	(2,906)	(3,693)	(4,031)	(3,975)	(4,144)
Less: Amortization of acquired intangible assets	(2,853)	(3,329)	(3,337)	(3,653)	(5,961)
Non-GAAP cost of revenue	\$109,817	\$121,655	\$138,948	\$154,328	\$174,053
GAAP gross profit	\$363,511	\$383,639	\$415,711	\$452,551	\$455,597
Add: Stock-based compensation expense & related payroll taxes	2,906	3,693	4,031	3,975	4,144
Add: Amortization of acquired intangible assets	2,853	3,329	3,337	3,653	5,961
Non-GAAP gross profit	\$369,270	\$390,661	\$423,079	\$460,179	\$465,702
GAAP gross margin	76%	75%	74%	74%	71%
Non-GAAP gross margin	77%	76%	75%	75%	73%
GAAP sales and marketing expense	\$214,011	\$219,359	\$236,309	\$251,138	\$271,600
Less: Stock-based compensation expense & related payroll taxes	(30,205)	(36,818)	(34,787)	(36,038)	(42,824)
Less: Amortization of acquired intangible assets	(388)	(417)	(362)	(659)	(1,250)
Non-GAAP sales and marketing expense	\$183,418	\$182,124	\$201,160	\$214,441	\$227,526
GAAP S&M expense as a % of revenue	45%	43%	42%	41%	42%
Non-GAAP S&M expense as a % of revenue	38%	36%	36%	35%	36%
GAAP research and development expense	\$115,089	\$134,557	\$120,956	\$141,887	\$150,972
Less: Stock-based compensation expense & related payroll taxes	(38,269)	(50,956)	(38,450)	(46,947)	(49,501)
Non-GAAP research and development expense	\$76,820	\$83,601	\$82,506	\$94,940	\$101,471
GAAP R&D expense as a % of revenue	24%	26%	22%	23%	24%
Non-GAAP R&D expense as a % of revenue	16%	16%	15%	15%	16%
GAAP general and administrative expense	\$87,658	\$96,987	\$95,906	\$108,760	\$95,019
Less: Stock-based compensation expense & related payroll taxes	(34,515)	(40,526)	(42,431)	(45,391)	(30,988)
Less: Acquisition-related and other expenses	(112)	-	-	(3,797)	(423)
Less: Lease impairment charges	-	(3,840)	-	(1,257)	-
Less: Legal reserve and settlements	-	-	-	2,886	-
Non-GAAP general and administrative expense	\$53,031	\$52,621	\$53,475	\$61,201	\$63,608
GAAP G&A expense as a % of revenue	18%	19%	17%	18%	15%
Non-GAAP G&A expense as a % of revenue	11%	10%	10%	10%	10%
GAAP loss from operations	(\$53,247)	(\$67,264)	(\$37,460)	(\$49,234)	(\$61,994)
Add: Stock-based compensation expense & related payroll taxes	105,895	131,993	119,699	132,351	127,457
Add: Amortization of acquired intangible assets	3,241	3,746	3,699	4,312	7,211
Add: Acquisition-related and other expenses	112	-	-	3,797	423
Add: Lease impairment charges	-	3,840	-	1,257	-
Add: Legal reserve and settlements	-	-	-	(2,886)	-
Non-GAAP income from operations	\$56,001	\$72,315	\$85,938	\$89,597	\$73,097
GAAP operating margin	(11%)	(13%)	(7%)	(8%)	(10%)
Non-GAAP operating margin	12%	14%	15%	15%	11%

Definitions & Methodologies

Paying Customers. We believe our ability to grow the number of paying customers on our network provides a key indicator of growth of our business and our future business opportunities. We define a paying customer at the end of the quarter as a person or entity who has generated revenue and has an active contract with us or one of our partners during such quarter, excluding (i) customers that were not acquired through ordinary sales channels, (ii) customers using only our registrar product, and (iii) customers using our consumer applications, such as 1.1.1.1 and Warp, which agreements and customers together represent an insignificant amount of our revenue. An entity is defined as a company, a government institution, a non-profit organization, or a distinct business unit of a large company. An active contract is defined as a customer relationship for which we have provided services during the quarter.

Paying Customers (> \$100,000 Annualized Revenue). While we continue to grow customers across all sizes, over time, our large customers have contributed an increasing share of our revenue. We view the number of customers with Annualized Revenue greater than \$100,000 as indicative of our penetration within large enterprise accounts. To measure Annualized Revenue at the end of a quarter, we take the sum of revenue for each customer in the quarter and multiply that amount by four. For example, if we signed a new customer that generated \$1,800 of revenue in a quarter, that customer would account for \$7,200 of Annualized Revenue for that year. Our Annualized Revenue calculation excludes (i) agreements that were not entered into through ordinary sales channels, (ii) revenue generated from customers using only our registrar product, and (iii) customers using our consumer applications, such as 1.1.1.1 and Warp, which agreements and customers together represent an insignificant amount of our revenue. Our Annualized Revenue metric also includes any usage charges by a customer during a period, which represents a small portion of our total revenue and may not be recurring. As a result, Annualized Revenue may be higher than actual revenue over the course of the year.

Dollar-Based Net Retention. Our ability to maintain long-term revenue growth and achieve profitability is dependent on our ability to retain and grow revenue generated from our existing paying customers. We believe that we will achieve these objectives by continuing to focus on customer loyalty and adding additional products and functionality to our network. Our dollar-based net retention rate is a key way we measure our performance in these areas. Dollar-based net retention measures our ability to retain and expand recurring revenue from existing customers. To calculate dollar-based net retention for a quarter, we compare the Annualized Revenue from paying customers four quarters prior to the Annualized Revenue from the same set of customers in the most recent quarter. Our dollar-based net retention includes expansion and is net of contraction and attrition, but excludes Annualized Revenue from new customers in the current period. Our dollar-based net retention excludes the benefit of free customers that upgrade to a paid subscription between the prior and current periods, even though this is an important source of incremental growth. We believe this provides a more meaningful representation of our ability to add incremental business from existing paying customers as they renew and expand their contracts.

Contracted Customers. Our contracted customers, which consist of customers that enter into contracts for our Enterprise subscription plan, have contracts that typically range from one to three years and are typically billed on a monthly or annual basis. Our agreements with contracted customers are tailored and priced to meet their varying needs and requirements. Enterprise subscription plan agreements for our contracted customers generally include a base subscription and a smaller portion based on usage or per seat.

Definitions & Methodologies

Annualized Recurring Revenue (ARR). Annualized Recurring Revenue, or ARR, is defined as the annualized value of our customer subscription contracts as of the measurement date, assuming any contract that expires during the next 12 months is renewed on its existing terms. ARR is calculated by multiplying the revenue in the fourth quarter of each respective year by four.

Annual Contract Value (ACV). Annual Contract Value, or ACV, is defined as the revenue that we would contractually expect to recognize over the term of the contract divided by the term of the contract in years.

Total Addressable Market. Total Addressable Market figures are Cloudflare estimates based on Gartner market forecasts and include the following Gartner sources: (1) Gartner, Forecast: Enterprise Network Equipment, Worldwide, 2024-2030, 1Q26, Vivek Tiwari, Christian Canales, Naresh Singh, 3/27/2026 (2) Gartner, Forecast Analysis: Edge Distribution Platforms, Worldwide, 2025, Brandon Medford, Jimmy Chuang, 1/26/2026 (3) Gartner, Forecast: Information Security, Worldwide, 2024-2030, 1Q26, Shailendra Upadhyay, Christian Canales, Ruggero Contu, Frank Marsala, Rahul Yadav, 3/26/2026 (4) Gartner, Forecast: IT Operations Management Software, Worldwide, 2022-2028, 2Q24 Update, Shailendra Upadhyay, Kanchi Bindal, Colin Fletcher, 10/16/2024 (5) Gartner, Forecast: Communications Services, Worldwide, 2024-2030, 1Q26, Grigory Betskov, Gaspar Valdivia, Khurram Shahzad, 3/26/2026 (6) Gartner, Forecast Analysis: Communications Platform as a Service, Worldwide, Lisa Uden-Farboud, Arunasree Cheparthi, 6/6/2025 (7) Gartner, Forecast: Enterprise Infrastructure Software, Worldwide, 2024-2030, 1Q26, Arunasree Cheparthi, Lisa Uden-Farboud, Robin Schumacher, Colin Fletcher, Sharat Menon, Kelli Smith, Saikat Ray, Nicholas Carter, Arthur Villa, Christian Canales, Brandon Medford, Amarendra ., Shailendra Upadhyay, 3/25/2026 (8) Gartner, Forecast: AI Spending, Worldwide, 2024-2029, 4Q25, Kay Arnott, Amarendra ., Adrian O'Connell, Hardeep Singh, Arunasree Cheparthi, Colleen Graham, Ranjit Atwal, Naresh Singh, Shailendra Upadhyay, Jon Erensen, Peter Middleton, 12/19/2025 (9) Gartner, Forecast Analysis: AI-Optimized IaaS, Hardeep Singh, Colleen Graham, 7/15/2025. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.