ANTHROP\C aws

Deploying secure production-ready agents with Claude Managed Agents on AWS

Jess Yan
Member of Technical Staff, Anthropic

Anton Aleksandrov
Pr. Solutions Architect, Agentic AI, AWS

Rohan Mehta
Solutions Architect, Serverless, AWS

Model capabilities have increased exponentially

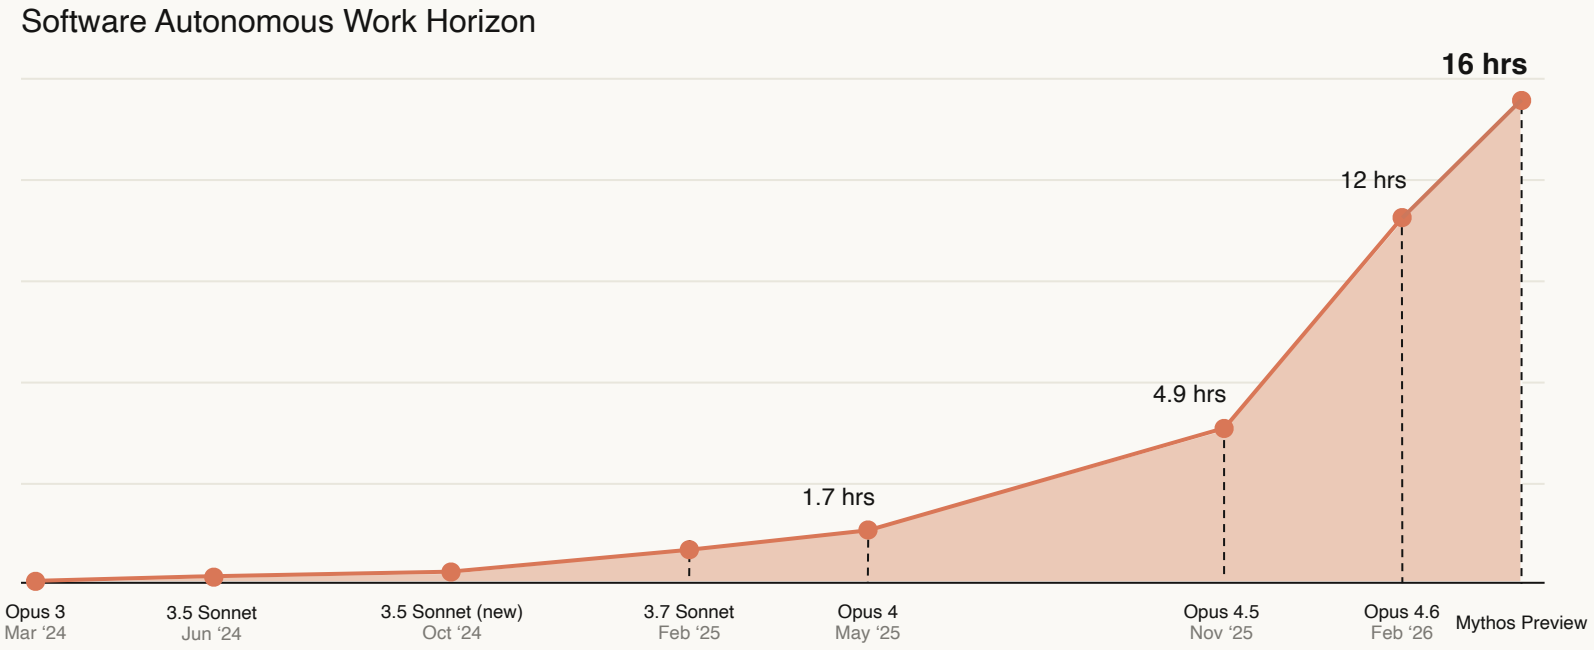

1 Chart depicts the task-completion time at which an AI agent is predicted to succeed with 50% reliability compared to the time it would take a human expert to complete the same task. METR, Task-Completion Time Horizons of Frontier AI Models (data as of May 2026).
2 Mythos Preview released in April 2026

Task complexity has increased alongside the models

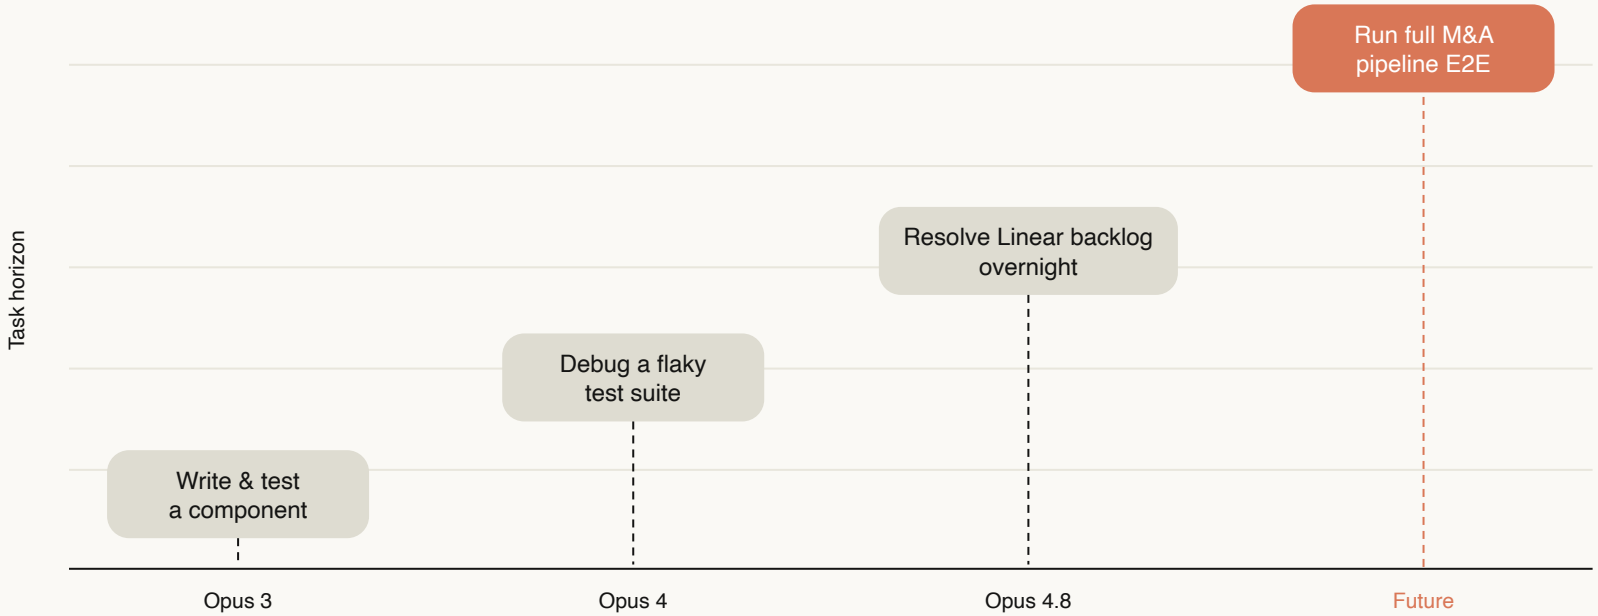

As agents operate on longer task horizons, we grant them more autonomy and access

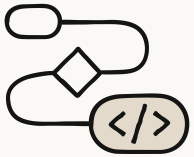

- Secure credentials
- Internal systems
- Proprietary code
- Identity and auth

We also expect more sophisticated interaction patterns

Conversational

Steer, queue messages, and interrupt

Outcome-oriented

Iterate until a satisfactory result is produced

Start and resume later

Pick up a task many days after kicking it off

Why teams build with Claude Managed Agents

Production-ready agent

The infrastructure to build and scale agents in your products. Credentials stay out of the sandbox, encryption is built in, and state persists automatically.

Fast iteration

Session tracing records what every agent did and why, with built-in analytics to improve performance.

The frontier

The only harness purpose-built for Claude. As the model improves, your agents improve with it.

Built for the frontier and the best agentic model

Model and harness built together: we test them together, so what we learn from one feeds back into both

Better performance: the harness is tuned to Claude specifically, not averaged across models.

Not every step needs the frontier model: Claude can plan the work and hand each step to a smaller model.

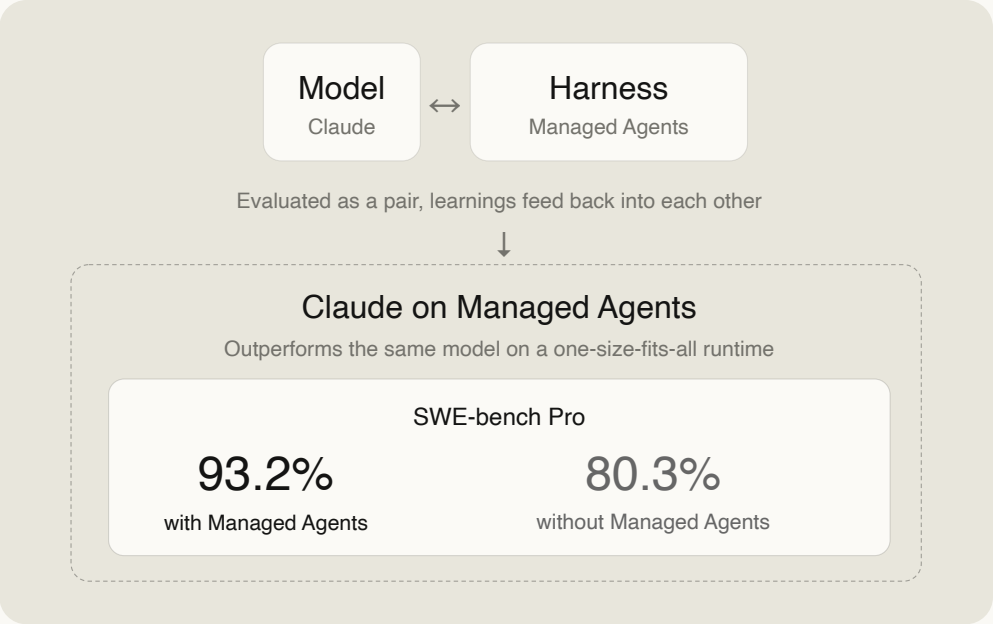

Built around three resources

Agent

Model · Prompt · Tools · MCP servers · Skills · Guardrails

Environment

Sandbox · cloud or self-hosted

Connected resources

Session

A running agent instance, working a task

- user message
- model call
- tool call
- tool result
- appended as it runs

Persisted · Resumable · Steer or interrupt mid-run

Every agent comes with frontier capabilities built in

Multiagent orchestration

Public beta

Claude delegates tasks to other agents, each with its own independent context window.

Self-hosted sandboxes

Public beta

Keep sensitive files, packages, and services in your own infrastructure — or use a managed provider.

Outcomes

Public beta

Claude iterates until it satisfies pre-defined exit criteria, grading its own work against a rubric.

Dreaming

Research preview

Between sessions, Claude reflects on past runs and codifies learnings into new memories.

Memory

Public beta

Claude reads and writes to memory stores, so every session gets progressively better.

MCP tunnels

Research preview

Reach MCP servers inside your private network without exposing them to the public internet.

Scheduled deployments

Public beta

Run agents on a schedule, so recurring work happens without having to trigger a run.

Credential vaults

Public beta

Securely manage OAuth tokens and environment variables for MCP and CLI/API authentication

Ship agents in your product, or run them inside your company

In your product

Agents your customers use

Notion

Notion built External Agents: Claude reads the connected pages and product docs, then hands its output back to the team for review. A single board can run thirty tasks at once.

Asana

Asana built AI Teammates: each teammate has its own account, so you can assign it a task or mention it in a comment. The teammate does the work, and a person makes the decisions.

Inside your company

Agents your employees use

Rakuten

Rakuten runs specialist agents across engineering, product, sales, marketing, and finance. Each one takes a week to deploy, and works inside Slack, Teams, and their own task system.

Other internal agents use cases

- Investigating production incidents and posting the root cause
- Triaging inbound tickets and drafting the first response
- Reviewing submissions against a compliance checklist
- Pulling a project's history into a status update for the team

Handle complex tasks with multiple agents

What it is
A lead agent breaks the job into pieces and delegates each one to a specialist with its own model, prompt, and tools.

How it works
Specialists work in parallel on a shared file system and contribute to the lead agent's overall context.

Developer control
Trace every step in the Claude Console, which agent did what, in what order, and why.

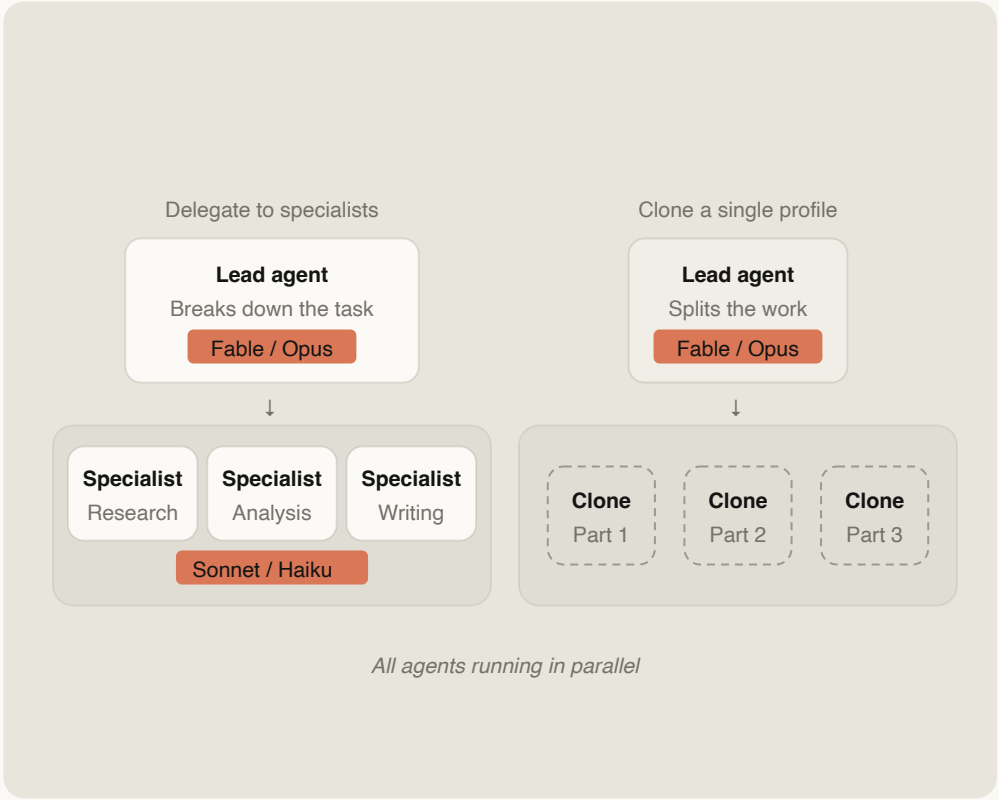

Deliver better outcomes

What it is
You write a rubric describing what success looks like, and the agent works toward it.

How it works
A separate grader evaluates the output against your criteria in its own context window.

Developer control
Define an outcome, let the agent run, and get notified by a webhook when it's done.

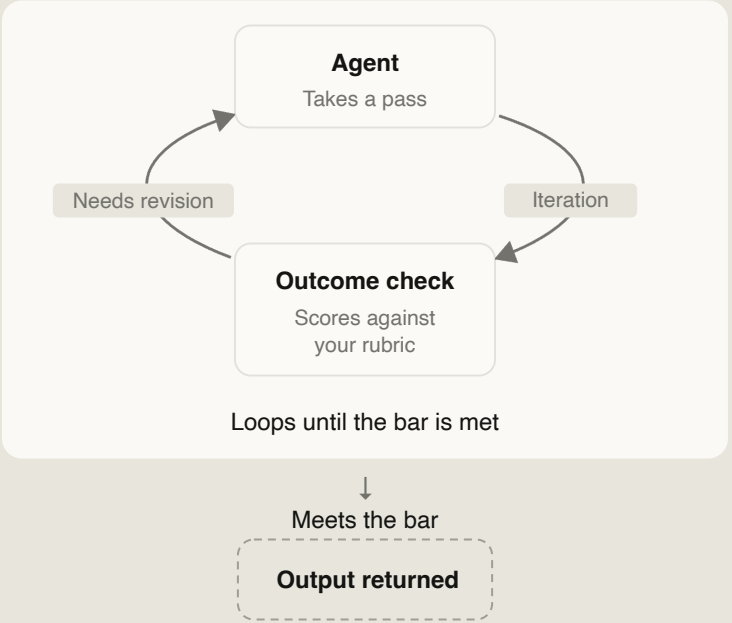

Memory enables learning during sessions

What it is
A filesystem that agents read and write to during sessions, so memories persist across runs.

How it works
Prior work, user preferences, and domain context are shareable across agents so they can learn from each other.

Developer control
You decide what each agent can access, with a full audit log of every change.

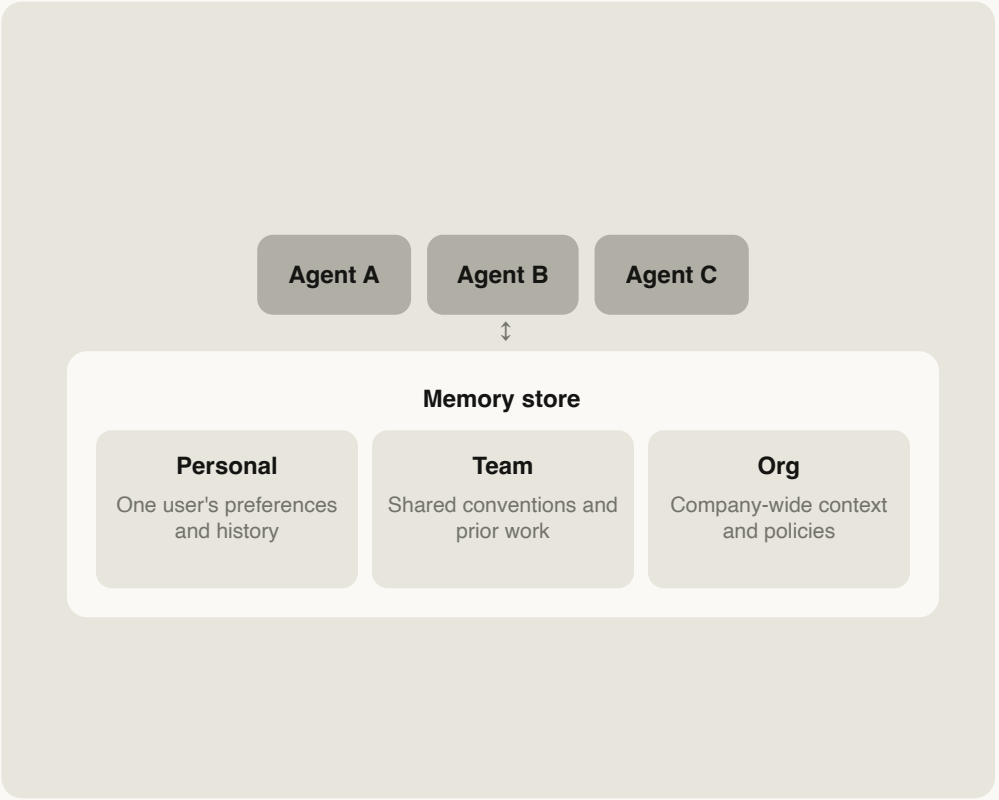

Dreaming enables learning between sessions

What it is
A scheduled process that reviews agent sessions, extracts patterns, and curates memories so agents learn over time.

How it works
Surfaces recurring mistakes, workflows, and preferences, and restructures memory to keep it high-signal over time.

Developer control
Updates memory automatically, or lets you review changes first, you decide how much autonomy the agent gets.

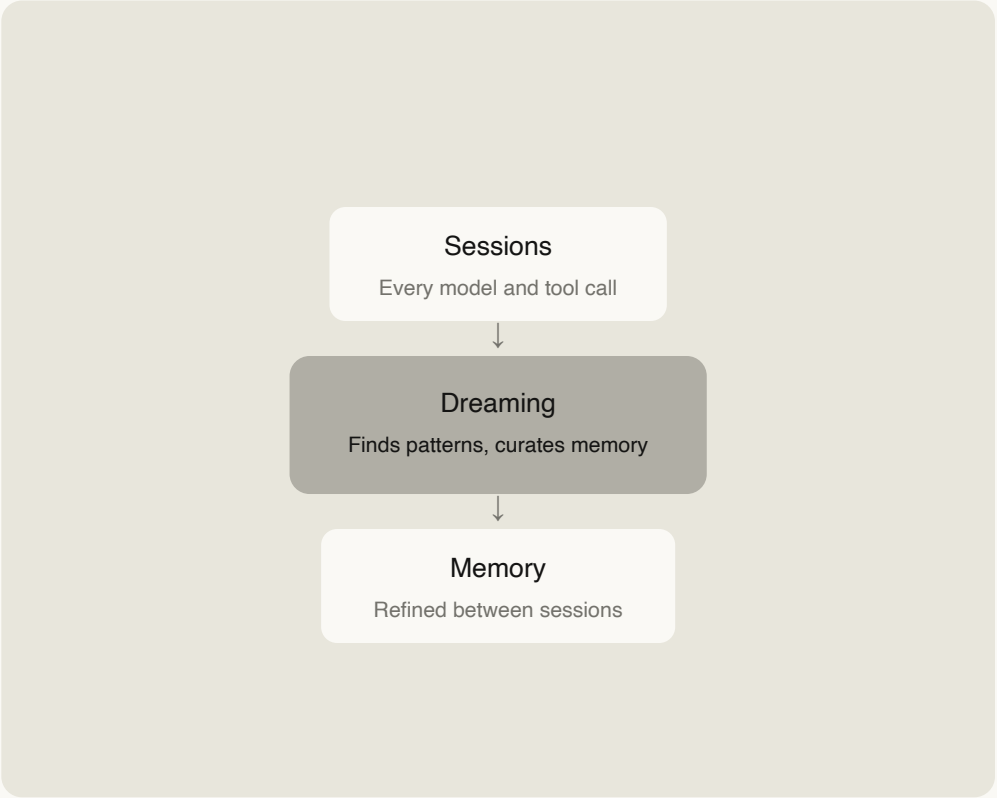

Scheduled deployments

Run agents on a schedule

What it is
Define a cron schedule and the agent runs on its own on a daily, weekly, or custom interval.

How it works
Memory and session history is shared between runs, so the agent knows what's already been done.

Developer control
Scheduling is built into Managed Agents, ready for monitoring, reporting, and upkeep work.

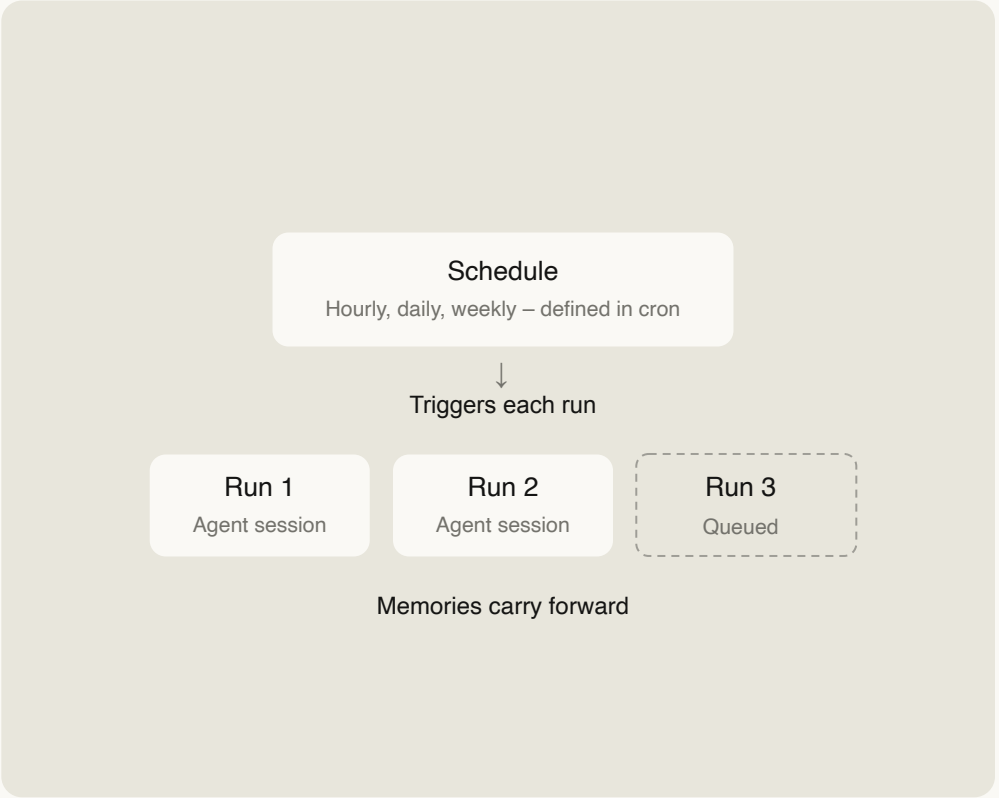

Using Lambda MicroVMs with Claude Managed Agents

Run untrusted code in a sandbox

Run untrusted code in a sandbox **you control**

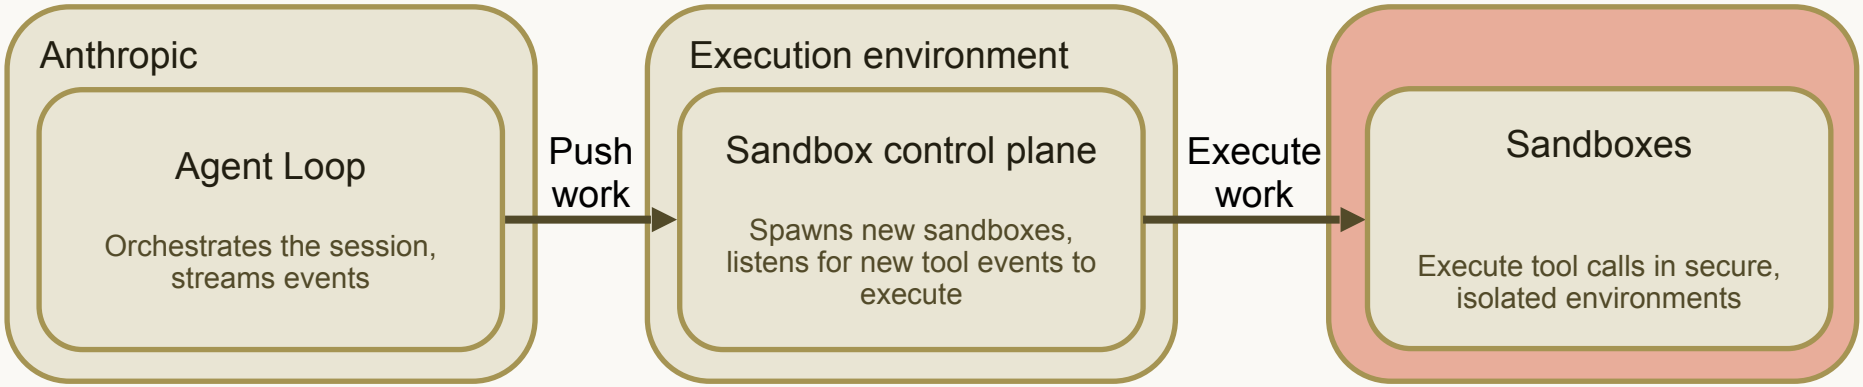

Run untrusted code in a sandbox **you control**

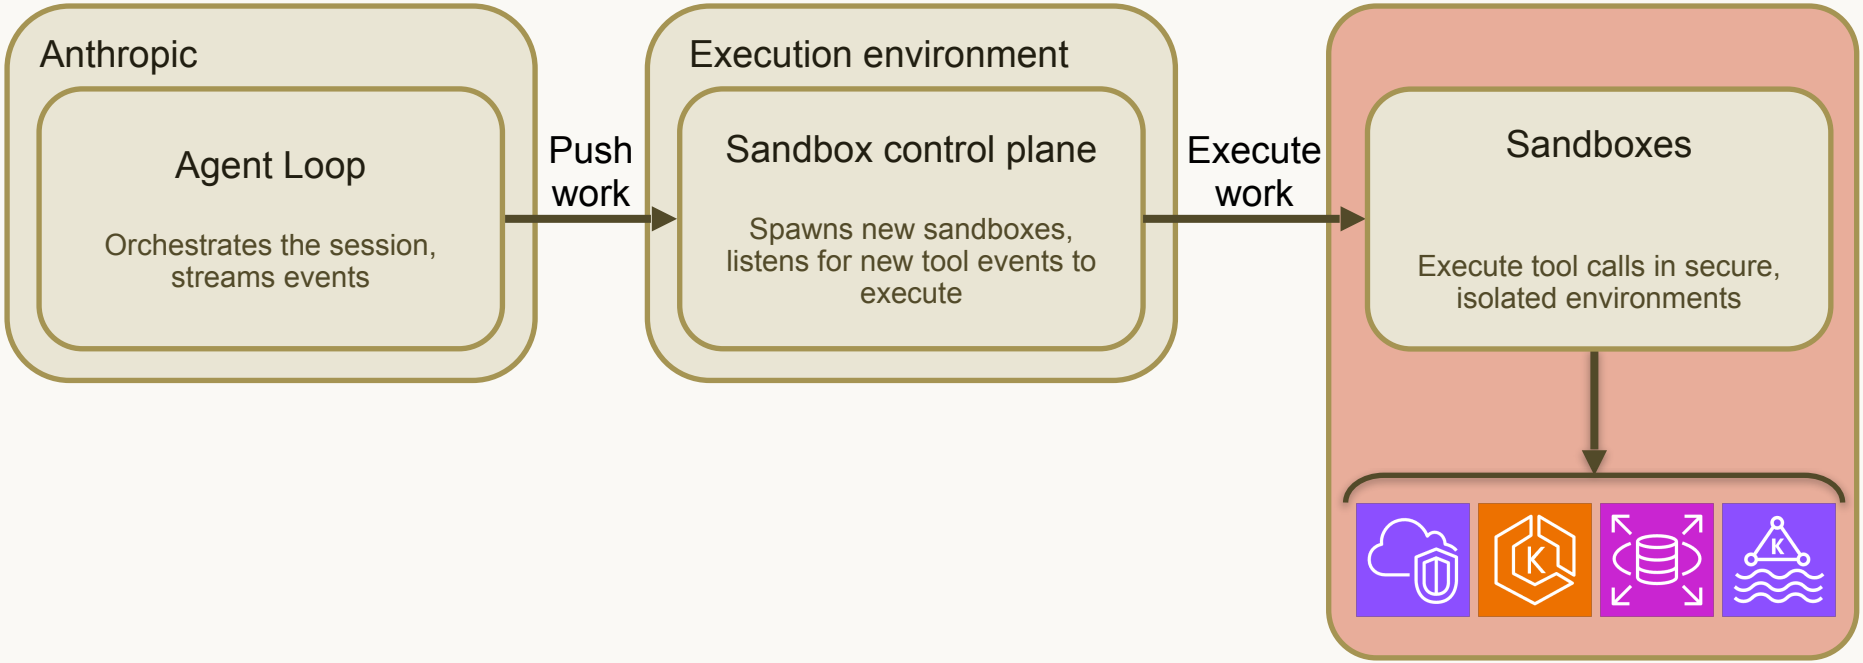

Agentic compute requirements in 2026

VM-level isolation

Container-level efficiency

Function-level speed and scale

Persistent compute state with suspend/resume capabilities

Familiar local/remote development experience

All-at-once, no infra management (Go Serverless!)

Typical compute trade-offs

	VM-level isolation	Near-instant startup	State persistence by design	No infra management	Suspend & resume
Virtual machines	✓	✗	✓	✗	✓
Containers	✗	⚠	✗	✗*	✗
Serverless functions	✓	✓	✗	✓	✗

Typical compute trade-offs

	VM-level isolation	Near-instant startup	State persistence by design	No infra management	Suspend & resume
Virtual machines	✓	✗	✓	✗	✓
Containers	✗	⚠	✗	✗*	✗
Serverless functions	✓	✓	✗	✓	✗
Lambda MicroVMs	✓	✓	✓	✓	✓

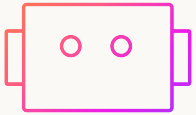

Lambda MicroVMs

Serverless compute for running user or AI-generated code in isolated, stateful execution environments.

Powered by ***Firecracker*** - the same tech behind Lambda's 15T+ monthly invocations.

Key capabilities

VM-level isolation

Snapshot-based near instant startup

Stateful suspend/ resume

Familiar Developer Experience

Core concepts

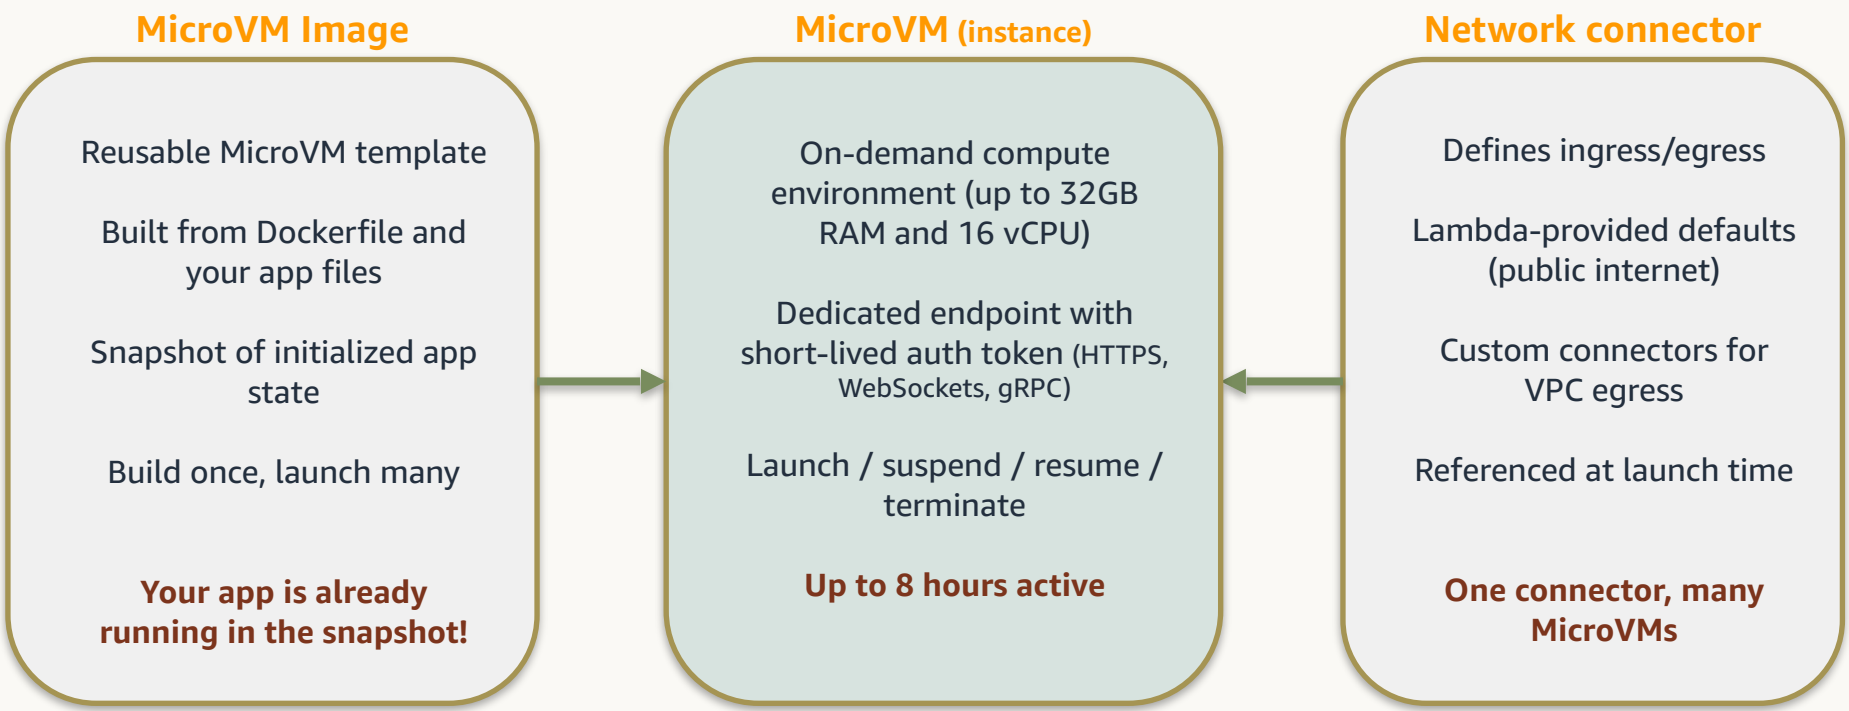

MicroVM lifecycle

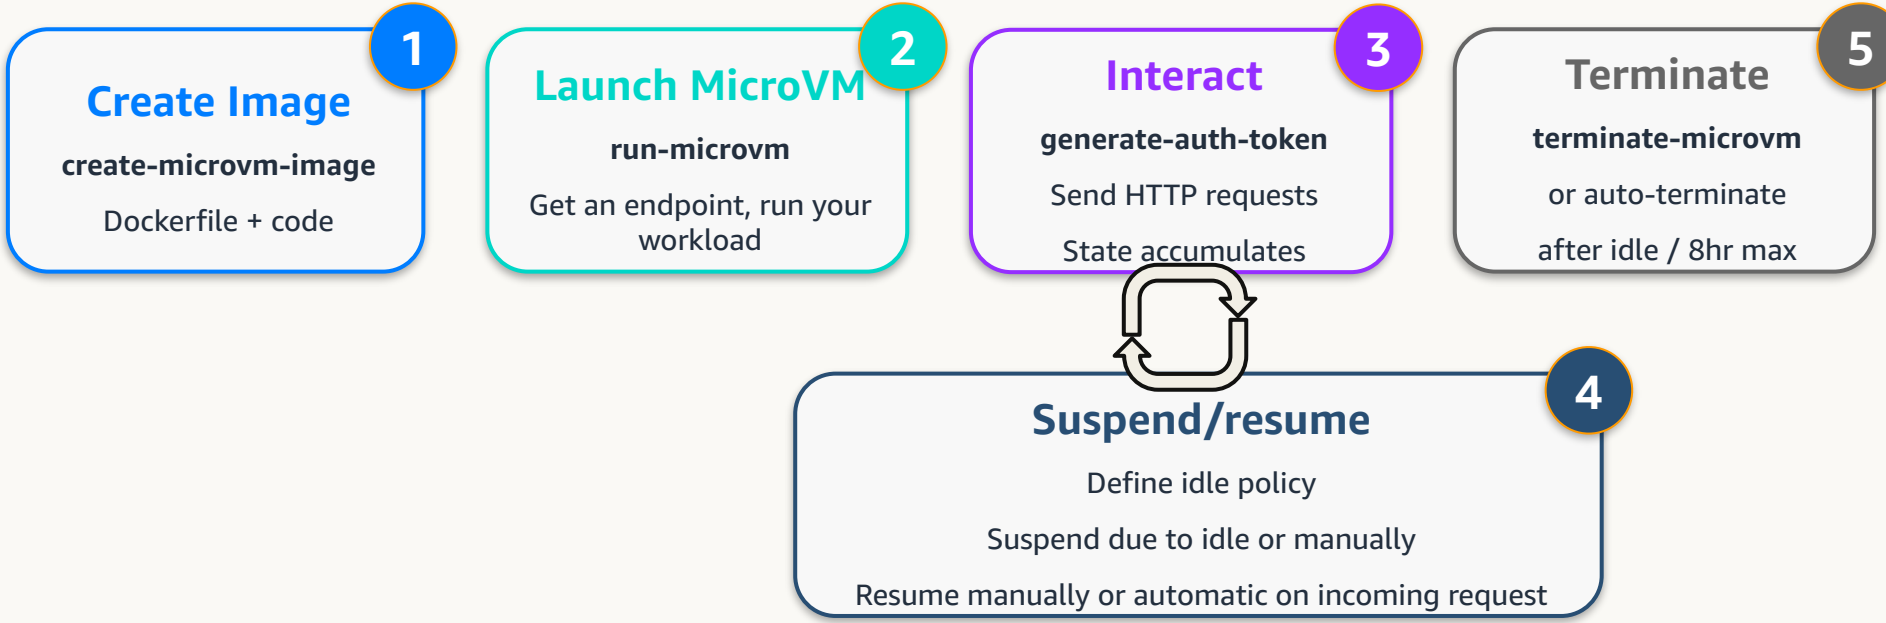

Claude Managed Agents with Lambda MicroVMs

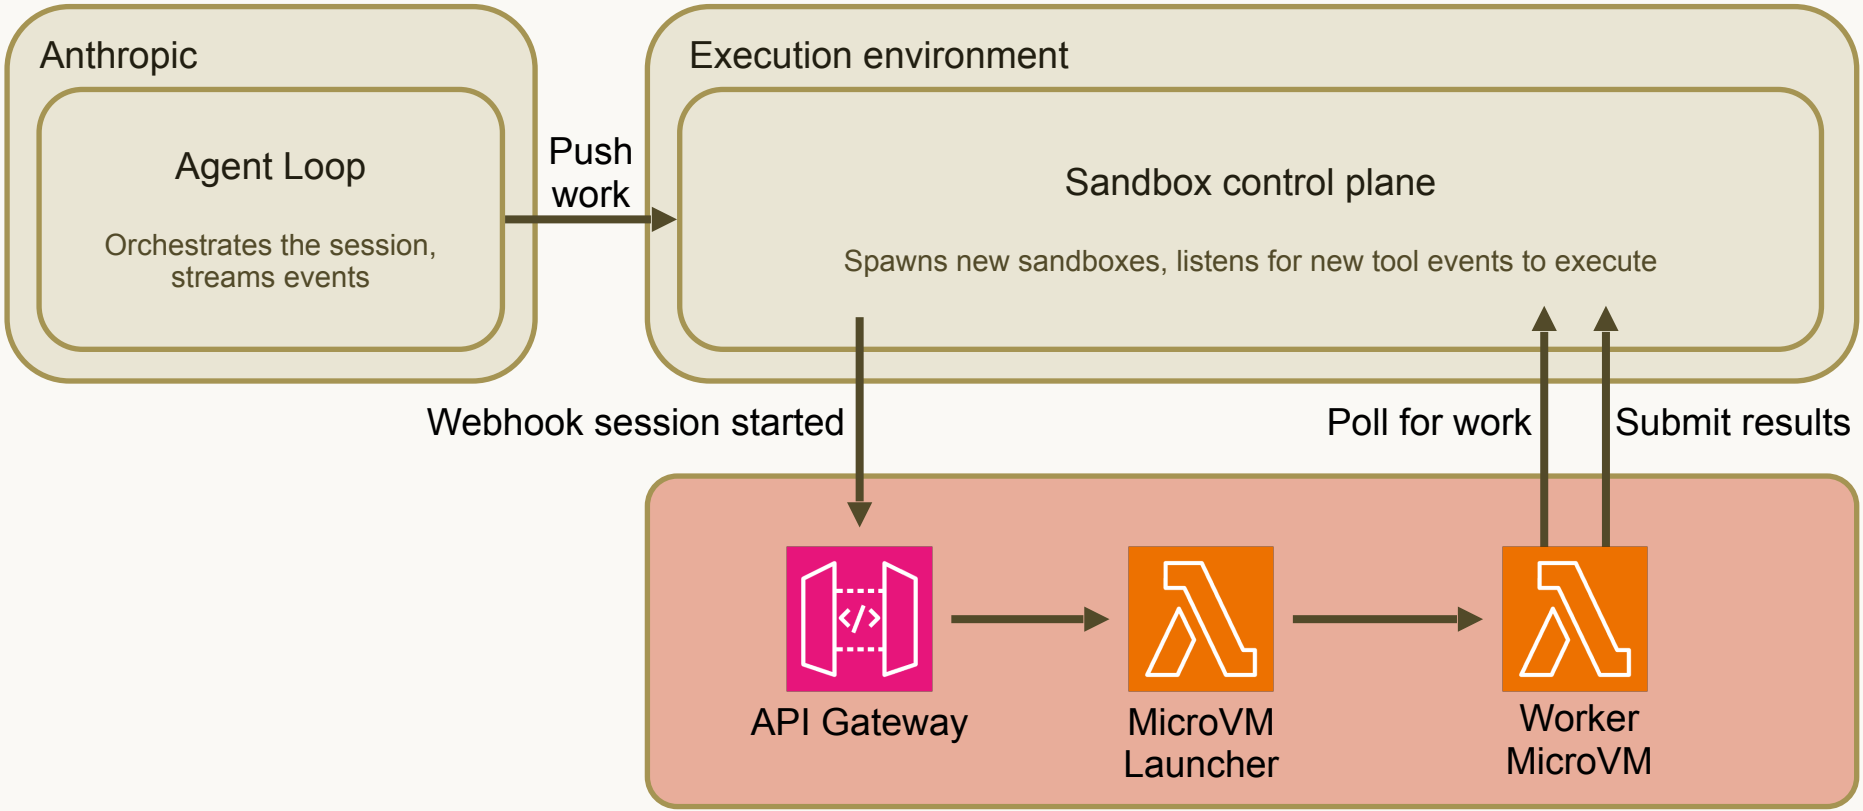

Demo time

Key takeaways

1

Managed agentic sandboxes with VM-level isolation and serverless speed

Firecracker gives you a real microVM boundary. Snapshot/restore gives you near-instant startups. You don't have to choose anymore.

2

State persists across requests and suspends

Your entire environment (memory, disk, processes) is persisted across interactions and idle periods. You pay storage, not compute, while suspended.

3

Give Claude Managed Agents access to your AWS domain

Anthropic runs managed agentic loop, Lambda MicroVMs provide secure, isolated sandboxes for running tools with access to your AWS domain

Continue your learning…

- Claude Managed Agents
- Claude Platform on AWS
- Lambda MicroVMs Dev Guide
- Using Lambda MicroVMs as a sandbox for Claude Managed Agents