Consultiply

• AI

MIP RESOURCE CENTER · MSP GUIDE · OCT 15, 2026

Trace the agent first. Widen its access second.

Anthropic opened OpenTelemetry tracing for Claude agents in public beta on Oct 14, 2026. What the traces show, what they do not, and how MSPs can turn trace review into a recurring service.

Claude

INDEPENDENT ANALYSIS OF ANTHROPIC'S ANNOUNCEMENT
NOT AFFILIATED WITH OR ENDORSED BY ANTHROPIC

What Anthropic announced

On October 14, 2026, Anthropic announced OpenTelemetry traces on the Claude Platform, in public beta. Agents built on Claude Managed Agents, the Messages API and the Claude SDKs can now send standard OpenTelemetry (OTel) traces to the observability tools a team already runs. Anthropic says "every agent turn, model call, and tool use arrives in your observability tool as a standard OpenTelemetry trace, with latency, tokens, and IDs."

Anthropic's launch post puts the operating model in two sentences: "Trusting an agent with more access starts with tracing its model calls, tool use, and MCP server calls. Review its traces, then widen its access as it proves reliable."

How it works

A trace is the full record of one request through a system; a span is one step in it, such as a model request or a tool use. According to Anthropic's OpenTelemetry docs, there are two ways to collect them:

- **API tracing.** Set up as a destination in the Claude Console, by someone with the Admin permission. It records each model run inside a request, Managed Agents turns and the server tools the API runs, for any client.
- **SDK tracing.** Set up in code, no special permission. It records client tools run by the SDK's tool runner, tools in a self-hosted sandbox and other requests such as file uploads.

Using both puts a request's SDK and API spans in one trace. For requests to Claude Platform on AWS, the API spans arrive in their own traces instead.

Spans follow the OpenTelemetry GenAI semantic conventions, so Claude data sits beside other providers in the same dashboards. Anthropic names 12 partner platforms (Arize AI, Braintrust, Databricks, Datadog, Grafana Cloud, Honeycomb, Langfuse, LangSmith, New Relic, PostHog, Sentry and Splunk) and says any tool that accepts a standard OTLP endpoint works.

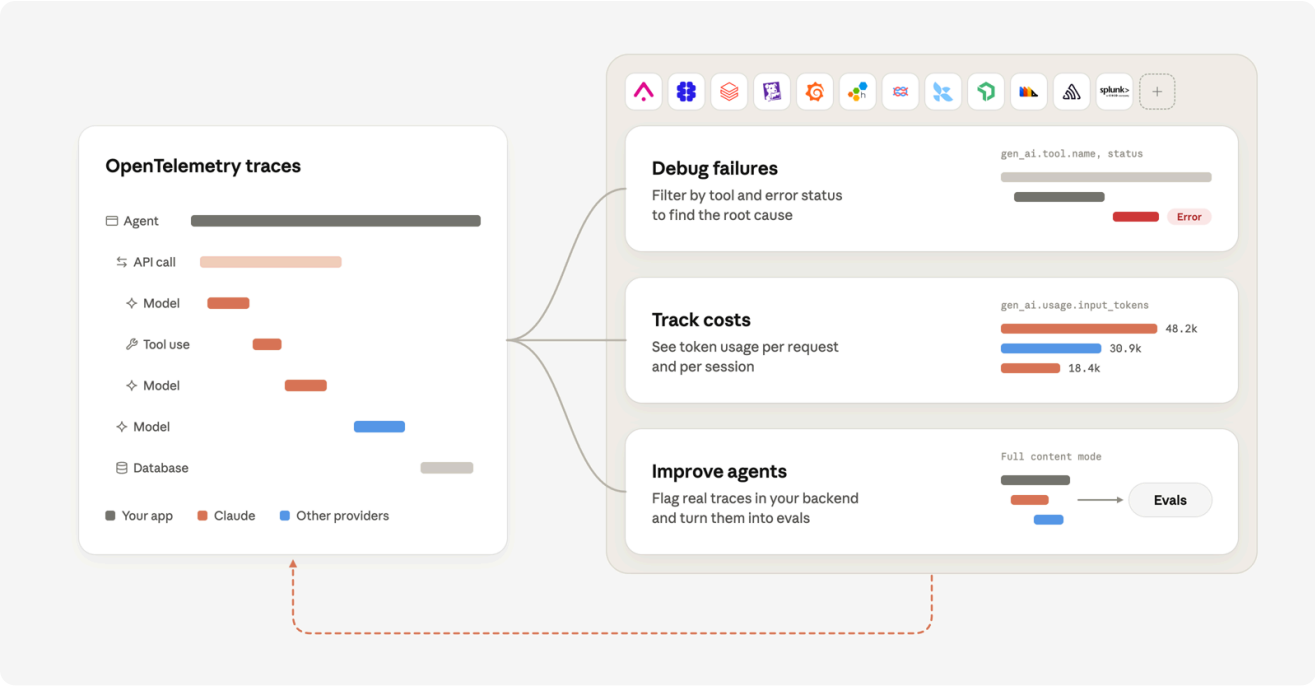


Image: Anthropic, Agent observability on the Claude Platform (Oct 14, 2026). Token counts are Anthropic's illustration, not measured data.

Two content modes, and why the choice matters

Each destination has a content mode. **Metadata only** is the default: identifiers, timing, status, token counts, and text you choose such as agent, tool and MCP server names. **Full content** adds model prompts and responses, tool inputs and tool results.

Full content is what makes a failed run readable. It also means client conversations are copied into the observability tool. Anthropic's docs are direct about it: "Your observability platform stores this content under its own access and retention settings, so check that they meet your organization's requirements." For organizations with a HIPAA, FERPA or K-12 data-handling restriction with Anthropic, destinations export only metadata, whatever mode is set.

Anthropic's numbers, and what they are

Few numbers come with this launch. Each one below is Anthropic's, or a customer's quoted by Anthropic. None was measured by Consultiply.

Claim	Measured	Status
3 surfaces covered: Managed Agents, Messages API, Claude SDKs	Not independently tested	vendor
12 named partner platforms, plus any OTLP endpoint	Not independently tested	vendor
Full content mode holds up to 0.75 MB of content per span by default	Not independently tested	vendor
Ramp says its own trace-scoring system has cut token usage by 10 to 20% on some agents and made sessions 30% faster	Customer claim about Ramp's system, not about the OTel feature itself	vendor

What it means for your clients

- **Agents become a monitored workload.** A client building on Claude can see agent steps beside its own logs and metrics, so a slow or failed step can be traced to the service it came from.
- **Trace data is client data.** In full content mode, prompts and tool results may contain customer records, so the observability tool's access and retention settings become part of the client's data boundary.
- **Token spend gets an owner.** Spans carry token counts per request and session, so spend can be attributed to an agent, team or feature instead of one monthly bill.
- **Regulated clients get metadata only.** Where a HIPAA, FERPA or K-12 restriction applies, plan dashboards that work without prompt content.

Offer design: what an MSP can package

These are Consultiply's suggestions for how an MSP could package its own services around the feature, not anything Anthropic offers.

- **Tracing setup as a scoped project.** Connect the destination, agree the content mode in writing, and build the first dashboards for failures, latency and token use.
- **Scheduled trace review.** A named reviewer reads a sample of traces on a fixed cadence and records what the agent did with its access.
- **Access recommendations.** Turn Anthropic's own rule into a decision log: which agent gets wider access, based on which traces, approved by whom.
- **Usage reconciliation.** Compare token totals from spans with the Usage API each month, since traces alone are not a complete record.
- **First response on agent incidents.** When an agent misbehaves, the trace is where the investigation starts. Write that into the client's incident process.

> Governance note: Full content mode copies prompts and responses into a third-party tool. Decide per client, in writing, which mode each destination uses and who can read the traces. Anthropic also notes that Claude Managed Agents isn't eligible for zero data retention, and that its retention stays the same when OpenTelemetry is on.

What to do this quarter

1. **List the clients building on Claude.** Note which use Managed Agents, the Messages API or the SDKs, and whether they already run an OTel-compatible tool.
2. **Pilot on your own agent first.** Trace an internal agent before offering it to a client, so your team learns to read the spans.
3. **Classify the data before choosing a mode.** Metadata only by default; full content only where the observability tool's retention meets the client's requirements.
4. **Write the widen-access criteria.** Decide what a run must show in its traces before an agent gets a new tool or permission.
5. **Reconcile monthly.** Check span token totals against the Usage API and investigate gaps.

Read with care

- **It is a beta.** Anthropic says OpenTelemetry export "is in beta" and provided under the Beta Services section of its terms. Expect setup and span details to change.
- **Traces can be incomplete.** In Anthropic's words: "Exported spans aren't guaranteed to be a complete record of your API usage. Their delivery is best effort, so a span can be lost." Anthropic points to the Usage API to track each API key's usage.
- **AWS requests trace differently.** For Claude Platform on AWS, API spans sit in separate traces from SDK spans.
- **Customer results are customers' claims.** The Ramp figures describe Ramp's own system, as quoted on Anthropic's page.
- **No price stated.** Neither the launch post nor the docs mention a separate charge for the feature. Observability platforms bill under their own terms.

> This guide is Consultiply's independent analysis of Anthropic's own announcement, translated for MSP operators. It is not affiliated with or endorsed by Anthropic, and it is not a product recommendation.

Read the full announcement: Agent observability on the Claude Platform.

Sources

Agent observability on the Claude Platform (Anthropic, Oct 14, 2026)

OpenTelemetry for the Claude API (Anthropic, Claude Platform docs, checked Oct 15, 2026)

ClaudeDevs launch post on X (Anthropic, Oct 14, 2026)

```
Read the full guide: consultiply.ai/mip-resource-center/claude-agent-tracing-opentelemetry-msp
```