Claude Managed Agents vs. Enterprise AI Platform Orchestration

Enterprise Security & Architecture

April 2026

Galent

Executive Summary

Claude Managed Agents (public beta, April 2026) is Anthropic's hosted agent runtime providing sandboxed execution, scoped permissions, and session tracing. Enterprise AI platforms similar to GalentAI Platform provides self-hosted, multi-agent orchestration with customer-controlled data residency and model portability. Both approaches are valid; the right choice depends on data sensitivity, compliance posture, and model-independence requirements.

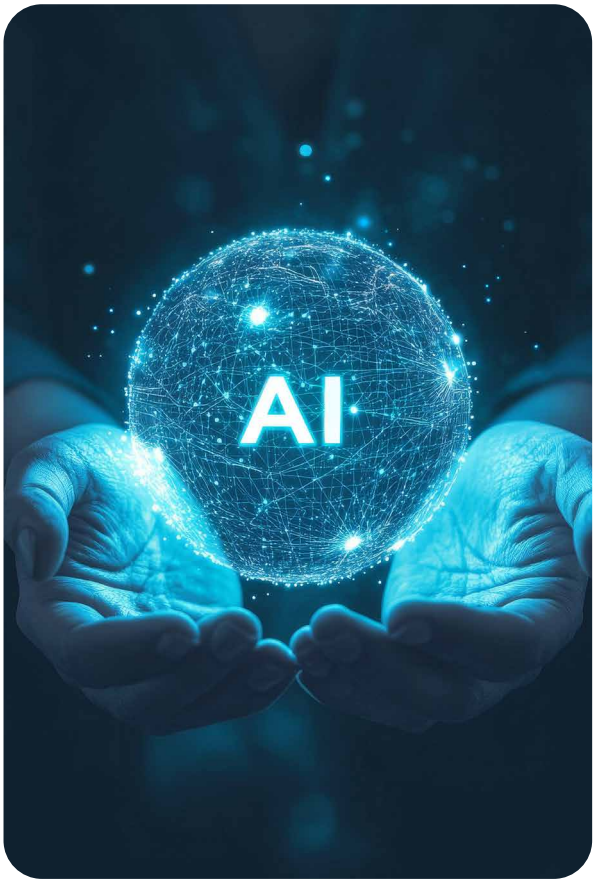

01

Infrastructure packaging, not a platform Multi-agent coordination and self-evaluation, the two capabilities that would make this genuinely powerful, are still not available.

02

Data residency is disqualifying for regulated industries Every tool call, every document, and every decision flows through Anthropic's cloud with no air-gap, no ZDR coverage, and no in-country data residency option, making it a non-starter for any client handling PHI, PII, financial records, or ITAR-controlled content.

03

Consumption pricing does not work at enterprise scale A 20-agent deployment running 8 active hours daily costs approximately $5,300 per month, scales linearly, carries no Batch API discount, and cannot be committed to as a flat-rate until a custom enterprise deal is negotiated.

04

Claude-only is a strategic liability Locking enterprise workloads to a single model vendor in a market where capabilities and pricing shift quarterly is a risk that surfaces in contract renewals and architecture debt, not in the original design review.

05

Session tracing is not enterprise governance Scoped permissions and session tracing are a starting point, not RBAC, not multi-tenant isolation, not per-agent cost attribution, and not the audit trail a CIO needs when answering to a board on AI risk.

Claude Managed Agents	Dimension	Enterprise AI Platform
Anthropic cloud (managed)	Hosting	Customer VPC / private cloud
Flows through Anthropic infra	Data Residency	Fully within customer boundary
Not covered by ZDR	Zero Data Retention	Customer-controlled retention
Claude models only	Model Choice	Multi-LLM via AI Gateway
Scoped permissions, session tracing	Governance	Full audit, tenant isolation, RBAC
Months (managed)	Time to enable cross team	Days (pre-built accelerators)
Per Developer and scale uncontrollably	Pricing	Reusable knowledge, context and Epic. Significantly optimized

When to Recommend Enterprise AI Platforms like GalentAI

Regulated data:
PHI, PII, financial records, proprietary source code, or ITAR/export-controlled content.

Enterprises requiring ZDR, air-gapped deployment, or in-country data residency.

Clients needing multi-LLM routing (Claude, GPT, Gemini, open-source) to avoid vendor lock-in.

Complex multi-agent orchestration with domain-specific context graphs and long-running workflows.

Organizations requiring centralized admin governance, multi-tenant isolation, and custom audit trails.

Predictable, flat-rate pricing aligned to enterprise procurement (no per-session or per-search variable costs).

Pre-built accelerators for App Modernization, Agent Builder, and Integration — measured time-to-value in days.

Deep domain customization:
proprietary prompts, fine-tuned retrieval, and client-specific knowledge graphs remain inside the customer boundary.

Unified observability:
LLM traces, token usage, cost attribution, and agent decision logs available natively without third-party APM.

Model independence:
swap or upgrade underlying models without rewriting orchestration — future-proof against vendor pricing or policy changes.

Enterprise support and co-engineering:
FDE delivery team partners directly on rollout, change management, and production hardening.

Proven across large enterprise engagements:
modernization, re-platforming, and AI implementation delivered end-to-end.

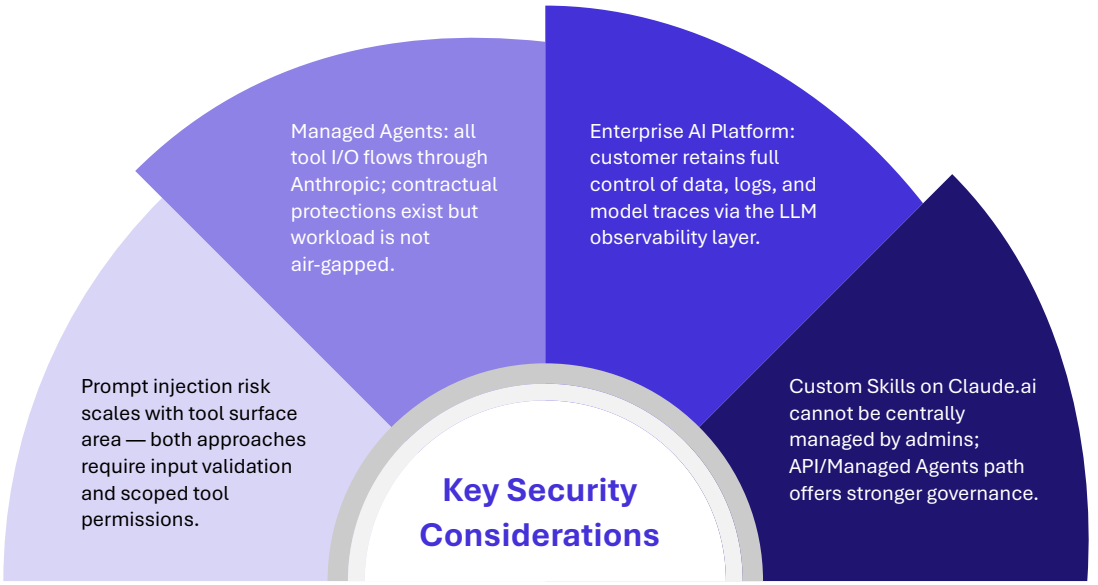

Recommendation

1

Use enterprise AI Platforms as the default for enterprise engagements

2

Offer Claude Managed Agents as a complementary accelerator for specific product-embedded agent use cases.

Thank You!

Galent