

The AI Agent Transparency Gap

What 15,220 marketplace listings reveal about what enterprises know — and don't know — about the AI agents they buy

October 2026 · Data from the SettleTop Global AI Registry · 6-minute read

AI agents have become remarkably easy to discover and buy. AWS, Microsoft, Google Cloud, Salesforce and ServiceNow now present agents as catalog products that enterprises can search, compare, subscribe to and deploy. The information needed to evaluate those agents has not kept pace with the buying experience.

SettleTop's Global AI Registry tracks **15,220 enterprise AI agent listings from 6,279 vendors across six marketplaces**. On every marketplace where the registry publishes an average, the mean provenance score sits between **2.0 and 2.8 out of 10**. Across the whole registry, **99.6% of listings fall in the Low provenance band**. Listings say what an agent does and how to buy it, but rarely the model, tools, data sources, data residency or permission scope.

THE CORE FINDING

Enterprise buyers are being told to evaluate security, integration, data handling, permissions, compliance and governance. Yet marketplace listings rarely disclose enough evidence to assess those criteria before purchase.

Key figures

15,220

AI agent listings tracked across six enterprise marketplaces

6,279

vendors and groups — about 2.4 listings per vendor

2.0–2.8

mean provenance score (out of 10) on every marketplace that publishes one

99.6%

of listings score Low provenance (15,153 of 15,220)

Source: SettleTop Global AI Registry, figures as of Sept. 29, 2026.

The marketplace is becoming part of procurement

Cloud marketplaces are already a major channel for enterprise software. Omdia forecasts hyperscaler marketplace sales rising from **\$30 billion in 2024 to \$163 billion by 2030**, while enterprises hold roughly **\$470 billion in cloud commitments** that can often be spent through them.

AWS Marketplace now lets buyers filter agents by deployment option, pricing, industry, compliance certification and integration protocol, and since November 2025 has offered agent mode, a conversational experience for discovering and comparing software.

What enterprises are told to evaluate


AWS advises buyers to compare features, pricing and compliance certifications, review integration requirements, and confirm agents meet security and regulatory needs. In procurement terms, that implies the evidence below.

Buyer criterion	Why it matters for an agent	Marketplace evidence needed
Capability / business fit	Can the agent accomplish the required task?	Use cases, functions, limitations, demonstrations
Integration	Can it work with the enterprise stack?	APIs, MCP/A2A support, systems supported, deployment model
Security & data handling	What enterprise data does it touch, and where does it go?	Data sources, residency, retention, authentication
Permissions & authority	What can the agent read, write or execute?	Scopes, roles, least-privilege options, approval boundaries
Compliance	Does the deployment meet organizational obligations?	SOC 2, ISO 27001, HIPAA, FedRAMP or other applicable evidence
Governance & traceability	Can the enterprise identify, monitor and audit the agent?	Owner, version, dependencies, change history, evidence trail
Economics	What will deployment and ongoing usage cost?	Price, consumption model, contract terms

Buyer criteria synthesized from AWS Marketplace buyer guidance and current enterprise governance guidance. SettleTop's registry measures whether these attributes are publicly disclosed — not whether an agent is intrinsically safe or high quality.

Every scored marketplace averages Low provenance

On each of the four marketplaces where SettleTop publishes an average, agents score between 2.0 and 2.8 out of 10. None reaches 3.0, where the Moderate band begins. Registry-wide, only 67 of 15,220 listings reach Moderate, and 98.8% are classified as High evidence risk.



Mean agent provenance score by marketplace, out of 10. Salesforce AgentExchange and the DRAI Agentic-AI Marketplace do not show a public average. Source: SettleTop Global AI Registry, Sept. 29, 2026.

Marketplace	Agent listings	Vendors	Mean provenance
Microsoft Marketplace	7,331	3,748	2.1
AWS Marketplace	5,309	2,422	2.0
Google Cloud Marketplace	1,908	302	2.2
Salesforce AgentExchange	541	365	Not published
ServiceNow Store	296	51	2.8
DRAI Agentic-AI Marketplace	23	1	Not published
Registry total	15,220	6,279*	—


*Vendors listing on more than one marketplace are counted once in the registry total, so marketplace vendor counts do not sum to 6,279. Source: SettleTop Global AI Registry marketplace pages, Sept. 29, 2026.

LOW PROVENANCE DOES NOT MEAN HIGH RISK

A low score means the public evidence is incomplete or hard to verify — not that the agent is insecure or poor quality. A well-engineered agent may disclose little; a transparent one can still be risky if it has broad permissions or a high-impact use case. Provenance measures evidence; risk requires a separate assessment of access, autonomy, controls and impact.


The price is on the label. The model often isn’t.

Consider **AI Refinery™**, Accenture’s agent platform listed on AWS Marketplace at a published **\$405,000 per year**. The listing identifies the vendor, delivery model, price and support path, and says customers can “flex between models” — but it does not name the model behind a deployed agent. The passport records model, framework, tools/MCP, data sources, data residency and permission scope as Unknown. **Six of twelve layers are disclosed.**



What the AI Refinery listing discloses across SettleTop's twelve passport layers. Overall provenance: 2.6 out of 10 (Low). Source: SettleTop agent passport, Sept. 2026.

The sub-scores show where the gap sits. Accenture itself scores **7.0 on vendor provenance** — its corporate identity is public and verifiable. But the agent scores **1.4 on build**, **1.4 on data** and **0.0 on compliance**.



AI Refinery provenance sub-scores, out of 10. Source: SettleTop agent passport for AWS Marketplace listing prodview-vyqjzsbqejgxy.

The brand-trust paradox

A trusted vendor and a transparent agent are not the same thing. Established firms score well on vendor provenance because their corporate records are public; their individual listings often disclose far less.

Accenture is the registry's largest publisher, with 535 listings; Deloitte has 269, Capgemini 233 and PwC 185. In SettleTop's September review of 1,158 listings from those four firms, **none linked a public code repository**. Commercial software commonly keeps code private, so buyers need other evidence: model and dependency disclosure, attestations, permission documentation, change history or independent verification.

THE PROCUREMENT SHIFT

Vendor due diligence asks: **"Do we trust this company?"** Agent due diligence adds a different question: **"What authority should we give this specific agent?"** The same vendor can supply one agent that summarizes public information and another that modifies records or executes transactions. Their risk profiles are fundamentally different.

IBM draws the same distinction in its September 22, 2026 guidance on third-party AI agents: governance “should reflect an agent’s autonomy, access, authority, and potential impact,” and “a trusted vendor relationship should not confer unrestricted authority on every agent.”

Why the transparency gap matters

Agents are not merely information interfaces. They call tools, access data, hold credentials and take actions — so missing information about data handling, permissions and dependencies is operationally important.

The visibility problem is already showing up inside enterprises. A Cloud Security Alliance survey of 418 IT and security professionals found that **82% had discovered previously unknown AI agents** in their environments in the past year and **65% had experienced an agent-related incident** in the prior 12 months. A separate Enterprise Management Associates study of 202 enterprise IT and security leaders found that **94% were confident their agents had no more access than needed**, yet **only 33% actually provisioned agents with least privilege**.



Perception versus reality in enterprise agent visibility and access. Sources: CSA / Token Security (Apr. 2026); EMA / Cequence Security (Aug. 2026). Both surveys cover agents from all sources, not only marketplace agents.

THE AI AGENT TRANSPARENCY GAP

Discovery has improved faster than verification. The gap is not simply that listings are incomplete; it is that enterprise buyers are being advised to evaluate attributes — security, data handling, permissions, compliance, integration and governance — that are not consistently visible at the point of discovery and purchase.

Seven questions to ask before you click Subscribe

1. Which model runs this agent, and can the model or configuration change without notice?
2. Which tools, APIs and MCP/A2A services can it call, and what can each one read, write or execute?
3. Where is enterprise data processed and stored, how long is it retained, and can it be used for training?
4. What permissions does the agent require, and is a least-privilege configuration available?

5. How much autonomy and authority does the agent receive, and which consequential actions require human approval?
6. Which attestations or controls apply to this specific agent or deployment — for example SOC 2, ISO 27001, HIPAA, FedRAMP or relevant NIST guidance?
7. Who built and maintains the agent, how are material changes disclosed, and which claims can be independently verified?

From discovery to verifiable identity

The agent marketplace has largely solved discovery before it has solved transparency. Low-disclosure agents are not necessarily bad agents — but they are harder to evaluate independently. Enterprises need to separate three questions: What can this agent do? What evidence do we have about what it is? And what risk does it create in our environment?

As third-party agents gain access to enterprise systems, every agent will need a verifiable identity and evidence record that follows it from discovery through deployment. The marketplaces have made agents easy to buy. The opportunity now is to make them just as easy to understand and verify.

EXPLORE THE REGISTRY

Search all 15,220 agent passports, compare marketplaces and vendors, and review the methodology at **registry.settletop.com** — free to browse, with deeper breakdowns available on a free account.

Media contact: James Carlson, Marketing · james@settletop.com

About the data

Registry-wide figures (listings, vendors, provenance and evidence-risk bands, marketplace averages and publisher counts) are taken from the public pages of the SettleTop Global AI Registry as of Sept. 29, 2026. The AI Refinery passport and the code-repository review of Accenture, Deloitte, Capgemini and PwC listings (1,158 listings) were captured Sept. 24–25, 2026; publisher totals have since grown to 1,222. Provenance is scored 0–10 across vendor, build, reach, data and compliance; scores below 3.0 are banded Low. “Unknown” means no public evidence was found at capture, not that the attribute does not exist.

Sources

- SettleTop Global AI Registry — overview, marketplaces, vendors and methodology. <https://registry.settletop.com>
- SettleTop agent passport — AI Refinery™ (Accenture, AWS Marketplace). <https://registry.settletop.com/agent/prodview-vyqizsbqeigxy>
- AWS Marketplace Buyer Guide, “Discovering AI agents and tools,” accessed Sept. 2026. <https://docs.aws.amazon.com/marketplace/latest/buyerguide/ai-agent-discovery.html>
- AWS, “AWS Marketplace introduces agent mode and AI-enhanced search,” Nov. 30, 2025. <https://aws.amazon.com/about-aws/whats-new/2025/11/aws-marketplace-agent-mode-ai-enhanced-search/>
- IBM, “How to effectively govern third-party AI agents across the enterprise,” Sept. 22, 2026. <https://www.ibm.com/think/perspectives/how-to-effectively-govern-third-party-ai-agents-across-the-enterprise>
- Cloud Security Alliance / Token Security, “Autonomous but Not Controlled,” Apr. 21, 2026. <https://cloudsecurityalliance.org/press-releases/2026/04/21/new-cloud-security-alliance-survey-reveals-82-of-enterprises-have-unknown-ai-agents-in-their-environments>
- Enterprise Management Associates / Cequence Security, “Agents Without Guardrails,” Aug. 31, 2026. <https://www.cequence.ai/news/ai-agent-governance-research/>
- Omdia, “Hyperscaler Cloud Marketplace Sales to Hit \$163 Billion by 2030,” Oct. 6, 2025. <https://www.businesswire.com/news/home/20251006925215/en/Omdia-Hyperscaler-Cloud-Marketplace-Sales-to-Hit-%24163-Billion-by-2030>

Media Contacts

For any media or questions regarding this report, please contact:

James Carlson
Marketing
james@settletop.com