SPECIAL REPORT

STATE OF AI AGENTS

Q2 2026

Market landscape, architecture patterns, enterprise adoption, and the growing security governance gap across the AI agent ecosystem.

$47B	144:1	92%
Market by 2030	NHI ratio	Blind to AI agents
MarketsandMarkets	Protego 2026	Saviynt

AUTHOR
Imiel Visser

PUBLISHED
April 2026

WEB
imiel.dev/explore

IMIEL.DEV / SPECIAL REPORT / Q2 2026

CONTENTS

01 EXECUTIVE SUMMARY

The AI agent ecosystem is undergoing a fundamental transformation. What began as experimental chatbot integrations has evolved into a sprawling landscape of autonomous systems that reason, plan, execute, and delegate across enterprise infrastructure.

$47B
Projected market by 2030
MarketsandMarkets

144:1
NHI to human identity ratio
Protego 2026

92%
Lack AI identity visibility
Saviynt/Cybersecurity Insiders

97%
Expect agent incident in 12mo
Security Boulevard

The central finding: enterprises are deploying AI agents at unprecedented speed while their security frameworks remain designed for human users. 97% expect a material incident within 12 months, yet only 6% of security budgets address the risk.

02 MARKET LANDSCAPE

The AI agent market is projected to grow from $5.1 billion in 2024 to over $47 billion by 2030, a compound annual growth rate of roughly 45%. Gartner projects 40% of enterprise applications will embed task-specific AI agents by 2026, up from less than 5% in 2025. McKinsey estimates agentic AI could automate tasks representing $3.5 to $4.1 trillion in economic value.

Source: Grand View Research 2024, MarketsandMarkets, Gartner, McKinsey

FOUNDATION MODEL PROVIDERS

Company	Platform	Category	Notable
Anthropic	Claude 4.x, Claude Code, MCP, Agent SDK	Foundation + Tools	MCP became de facto agent protocol
OpenAI	GPT-4.1, Codex, Operator, Agents SDK	Foundation + Tools	Acquired Windsurf (~$3B)
Google	Gemini 2.x, ADK, A2A Protocol, Agentspace	Foundation + Protocol	A2A for agent-to-agent interop
Meta	Llama 3.x, 4.x (open source)	Foundation (open)	Largest open-weight ecosystem
xAI	Grok 4.x	Foundation	Consumer-focused via X platform
Mistral AI	Mistral Large, Le Chat agents	Foundation (EU)	European AI leader, $640M Series B
Cohere	Command R+, enterprise RAG	Foundation (enterprise)	Data privacy focused deployments

PLATFORM AND ENTERPRISE PLAYERS

Company	Platform	Category	Notable
Microsoft	Copilot, AutoGen, Semantic Kernel, Azure AI Agent Service	Platform	70%+ Fortune 500 on Copilot
AWS	Bedrock Agents, Amazon Q Developer, Q Business	Cloud Platform	Multi-agent orchestration in Bedrock
Salesforce	Agentforce 2.0	Enterprise SaaS	Autonomous agents for sales/service/commerce
ServiceNow	Now Assist, Agent Workspace	Enterprise SaaS	IT workflow automation agents
SAP	Joule AI copilot, BTP agents	Enterprise ERP	Business process automation
Databricks	Mosaic AI Agent Framework	Data Platform	Agent evaluation and serving
Snowflake	Cortex Agents	Data Platform	Data analysis and retrieval agents

AGENT TOOLING AND INFRASTRUCTURE

Company	Product	Focus	Notable
LangChain	LangGraph, LangSmith	Orchestration + observability	95K+ GitHub stars, 100K+ LangSmith users
Cursor/Anysphere	Cursor IDE	AI-first code editor	$9B+ valuation, rapid growth
Cognition	Devin	Autonomous coding agent	$2B valuation, $175M Series A
Sierra AI	Customer experience agents	Enterprise CX	$4.5B valuation, Bret Taylor CEO
Harvey AI	Legal AI platform	Vertical (legal)	Multiple AmLaw 100 deployments
Perplexity	Answer engine + agents	Search + reasoning	$9B+ valuation
Replit	Replit Agent	Code generation + deploy	End-to-end app building
Vercel	v0 AI	Frontend generation	AI-powered UI development
Intercom	Fin AI Agent	Customer support	Autonomous resolution agent

Company	Product	Focus	Notable
Adept AI	ACT-2 model	Computer-use agent	Acquired by Amazon (2024)

FUNDING AND DEALS

Capital concentration reached historic levels in 2024-2025. The top six rounds alone totaled over $30 billion, with valuations reflecting expectations of an agent-driven platform shift.

Company	Amount	Valuation	Date
OpenAI	$6.6B round	$157B	Oct 2024
xAI	$6B round		Dec 2024
Anthropic	$10B+ cumulative		2024-2025
Windsurf/OpenAI	~$3B acquisition		Early 2025
Cursor/Anysphere	Series B+	$9B+	2025
SpaceX/Cursor	First-dibs deal option	$60B option	April 2026
Sierra AI	$175M Series B	$4.5B	Early 2025
Cognition (Devin)	$175M Series A	$2B	2024
Mistral AI	$640M Series B		June 2024
Perplexity AI	Multiple rounds	$9B+	2024-2025
CrewAI	$18M Series A		2024
LangChain	$25M Series A		2024
Astrix Security	$85M total		2025
Oasis Security	$75M		2025

Sources: TechCrunch, Bloomberg, The Information, company disclosures

PROTOCOL LANDSCAPE

Two open protocols emerged as standards for agent interoperability, creating the foundational communication layer for the agentic ecosystem.

Protocol	Scope	Creator	Adoption Status
MCP	Agent to Tool	Anthropic	De facto standard; adopted by OpenAI, Microsoft, Cursor, GitHub, thousands of servers
A2A	Agent to Agent	Google	Early adoption; Salesforce, SAP, Atlassian, ServiceNow as launch partners
OpenAPI	API description	OpenAPI Initiative	Foundation for tool definitions; complementary to MCP

03 ARCHITECTURE PATTERNS

Production architectures have stabilized around composing multiple patterns: Plan-and-Execute at the top level, ReAct within each step, hierarchical multi-agent for complex domains, and human-in-the-loop for safety-critical decisions. Pure autonomous agents remain experimental; the systems shipping at scale all include structured oversight.

PATTERNS IN PRODUCTION

Pattern	Maturity	Description	Shipped By
ReAct	HIGH	Alternates reasoning and tool-calling in a loop	OpenAI, Anthropic, LangChain
Plan-and-Execute	MED-HIGH	Planner generates full plan, executor runs steps	LangGraph, AutoGen
Hierarchical Multi-Agent	MED-HIGH	Supervisor delegates to specialist agents	Claude Agent SDK, CrewAI, OpenAI
Tool-Use (parallel)	HIGH	Multiple tool calls issued in a single turn	All major providers
Human-in-the-Loop	HIGH	Approval gates, escalation, collaborative editing	All enterprise deployments
RAG-Augmented	HIGH	Retrieval as a tool in the agent loop	LlamaIndex, LangChain
Agentic RAG	MEDIUM	Agent decides how/what to retrieve iteratively	LlamaIndex, GraphRAG
Swarm / Peer-to-Peer	LOW	Flat multi-agent coordination without supervisor	OpenAI Swarm (experimental)
Autonomous Long-Running	LOW	Fully autonomous multi-hour tasks	Cognition (Devin), SWE-agent

FRAMEWORK LANDSCAPE

The developer tooling ecosystem for building AI agents expanded rapidly. The table below covers the major frameworks, their community size, and primary use case focus.

Framework	Stars	Maintainer	Lang	Focus
LangChain	95K+	LangChain Inc	Py/JS	General-purpose LLM apps
LlamaIndex	38K+	LlamaIndex Inc	Py/JS	RAG + agentic retrieval
AutoGen	35K+	Microsoft	Py/.NET	Multi-agent conversation
CrewAI	25K+	CrewAI Inc	Python	Role-based multi-agent
Semantic Kernel	22K+	Microsoft	C#/Py/Java	Enterprise SDK, Azure native
LangGraph	10K+	LangChain Inc	Py/JS	Stateful orchestration
Haystack	18K+	deepset	Python	Pipeline-based RAG + agents
Pydantic AI	New	Pydantic team	Python	Type-safe agent framework
Claude Agent SDK	New	Anthropic	Python	Agentic loops + tool use
OpenAI Agents SDK	New	OpenAI	Python	Handoffs + guardrails
Google ADK	New	Google	Python	Vertex AI integration

GitHub star counts approximate as of Q1 2025. 'New' indicates launched in 2025.

OBSERVABILITY AND EVALUATION

Tool	Focus	Notable
LangSmith	Tracing, evaluation, monitoring	100K+ users, market leader in LangChain ecosystem
Langfuse	Open-source LLM observability	Strong OSS alternative to LangSmith
Braintrust	Evals, logging, prompt management	Dataset-driven evaluation with scoring
Helicone	LLM request logging, cost tracking	Open-source core, good cost visibility
Arize Phoenix	Open-source LLM observability	OpenTelemetry-based tracing
Weights and Biases Weave	LLM tracing and evaluation	Extension of W&B; ML platform

Tool	Focus	Notable
Inspect AI (UK AISI)	Agent safety evaluation	Open-source, safety benchmarks

KEY CHALLENGES

Context Window Management. Even with 128K-1M+ token windows, agents exhaust context over many steps. Solutions: summarization, selective context, hierarchical memory (working + episodic + semantic).

Tool Selection and Routing. Agents with 10-100+ tools struggle with selection. Retrieval-based tool selection, hierarchical organization, and fine-tuned routing models are active engineering areas.

Multi-Step Reliability. 95% per-step reliability = 60% over 10 steps. Plan-and-Execute with verification, self-reflection (Reflexion pattern), and structured workflows improve end-to-end reliability.

Cost Optimization. $0.10-$1.00+ per agent invocation. Model routing (cheap models for simple steps), prompt caching, structured outputs, and batch APIs at 50% discount all reduce costs.

Latency. Multi-step agents take 10-30+ seconds. Streaming, parallel tool execution, speculative execution, and async background processing are the primary mitigations.

04 ENTERPRISE ADOPTION

McKinsey found 72% of organizations adopted AI in at least one business function, with generative AI usage doubling in under a year. Gartner predicted 75% of enterprise developers would use AI coding assistants by 2028 (up from ~15% in 2024). The pilot-to-production gap is closing: 60-70% have pilots, 20-30% are in full production.

Source: McKinsey Global Survey on AI 2024, Deloitte Q1 2024, Gartner

ADOPTION BY INDUSTRY

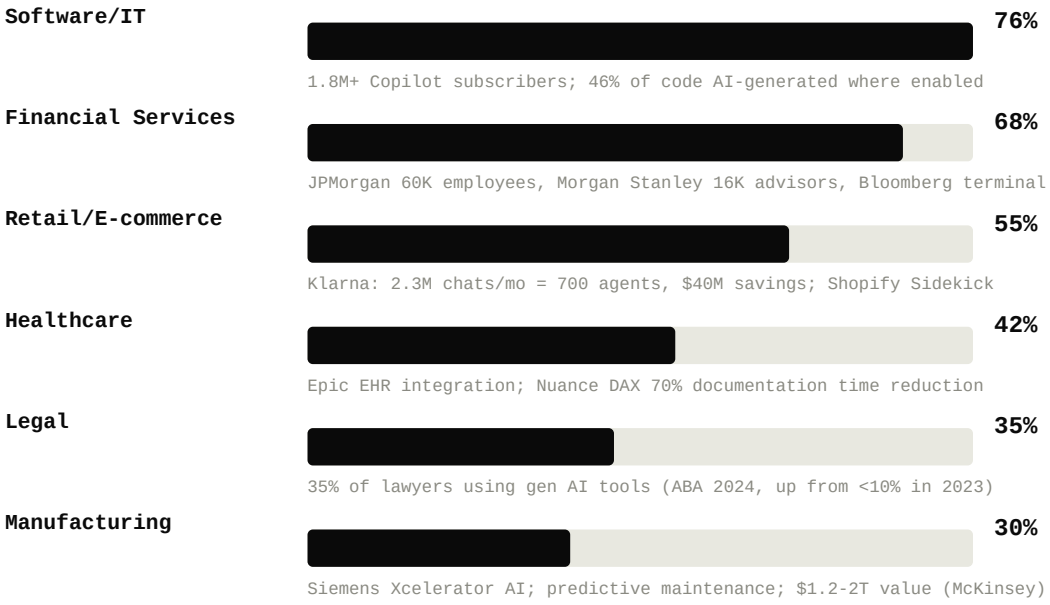

RETURN ON INVESTMENT

55%
Faster dev tasks
GitHub Copilot

40%
Higher quality
BCG/Harvard

$80B
Contact center savings
Gartner

Metric	Impact	Detail	Source
General productivity	20-70% boost	Depending on task complexity	McKinsey 2024
Consultant tasks	25% faster	With 40% higher quality output	BCG/Harvard 2023

Metric	Impact	Detail	Source
Support resolution	14% more/hr	Less experienced workers gain most	Nielsen Norman Group
Revenue increase	12.1% average	From gen AI deployments	Deloitte 2024
Cost reduction	11.5% average	From gen AI deployments	Deloitte 2024
Klarna savings	$40M/year	AI customer service automation	Klarna 2024

DEVELOPER ADOPTION

Metric	Value	Source
GitHub Copilot paid subscribers	1.8M+, 50K+ enterprise customers	GitHub early 2025
Code generated by Copilot	46% of code in files where enabled	GitHub
Developers using AI tools	76% using or planning to use	Stack Overflow 2024
Daily AI tool usage	50% of adopters use daily	JetBrains 2024
Cursor valuation	$9B+ (up from $400M in early 2024)	The Information
AI spending increase planned	8-12% average by CIOs	Gartner CIO Survey 2024
Worldwide AI spending 2026	$300B+ projected	IDC Spending Guide

> **"Klarna AI handled 2.3 million conversations in its first month, performing the work equivalent of 700 agents, with satisfaction scores on par with human agents."**
>
> Klarna, February 2024

BARRIERS TO ADOPTION

Barrier	Cited By	Source
Data privacy and security	55-70%	Gartner, Deloitte, McKinsey
Accuracy and hallucination risk	45-60%	Multiple surveys
Regulatory uncertainty	40-50%	Financial services, healthcare, EU
Lack of clear ROI measurement	35-45%	Deloitte 2024
Integration with legacy systems	30-45%	Deloitte 2024
Data quality and availability	35-40%	Deloitte 2024
Talent shortages	40%	McKinsey 2024

WORKFORCE IMPACT

Metric	Value	Source
Jobs displaced globally by 2027	83M (69M new roles created, net 14M)	WEF Future of Jobs 2024
Global jobs exposed to AI	40% (60% in advanced economies)	IMF Jan 2024
LinkedIn AI job posting growth	300%+ year-over-year	LinkedIn Economic Graph
Amazon upskilling investment	$1.2B for 300K employees	Amazon 2023
PwC AI training commitment	$1B over three years	PwC 2024
Orgs with AI reskilling programs	38% (up from 22%)	McKinsey 2024

05 SECURITY + GOVERNANCE

The most urgent finding: as AI agents proliferate, security teams are losing visibility. 92% of CISOs lack full visibility into AI identities. 97% expect a material AI-agent-driven security incident within 12 months. Yet only 6% of security budgets address the risk.

THE IDENTITY CRISIS

92%
Lack visibility into AI identities
Saviynt/Cybersecurity Insiders

144:1
NHI to human identity ratio
Protego 2026

97%
NHIs have excessive privileges
CSO Online

86%
No AI identity access policies
Saviynt 2026

75%
Found unsanctioned AI tools
Saviynt 2026

6%
Security budgets address AI risk
Bessemer

ADDITIONAL IDENTITY DATA

Metric	Value	Source
NHI growth year-over-year	44% (2024 to 2025)	Cybersecurity Tribe
Orgs with NHI breaches	50% already suffered a breach	Protego
Incidents involving machine IDs	68% of IT security incidents	Protego
Credential rotation compliance	92% fail 90-day rotation	SANS 2026
Orgs with real-time agent inventory	Only 21%	Strata.io
Formal AI identity lifecycle policy	Only 22% (78% have none)	IANS Research
Shadow AI prevalence	98% of orgs have unsanctioned AI	Nutanix

Metric	Value	Source
Shadow AI breach cost premium	+$670K per incident ($4.2M avg)	Second Talent

NOTABLE INCIDENTS

Incident	Date	Type	Impact
Vercel / Context AI	Apr 2026	OAuth supply chain	Hundreds of orgs; $2M BreachForums listing
CrowdStrike GenAI	2025-2026	Prompt injection	90+ organizations exploited via GenAI tools
OpenClaw Agent Crisis	2026	Malicious skills	21K exposed instances; 1,184 malicious skills
Drift/Salesforce OAuth	2025	OAuth token theft	700+ organizations via single integration
GeminiJack (Google)	Q4 2025	Zero-click exploit	Shared Doc/Calendar triggered data exfil
Reprompt (MS Copilot)	Q4 2025	Session hijacking	Injected prompts hijacked user sessions
MCP Tool Poisoning	2025-2026	Supply chain	492 MCP servers with zero authentication
Claude Code RCE	2026	Config poisoning	Remote code execution via repo config files

Sources: TechCrunch, BleepingComputer, CrowdStrike, Cycode, Reco AI, Lakera, Trend Micro

THREAT LANDSCAPE

340%

Injection increase YoY

Wiz Research

89%

Identity weaknesses

Unit 42 2026

72min

Access to exfiltration

Unit 42

Statistic	Detail	Source
73%	Production AI deployments hit by prompt injection	Markaicode
48%	Identify agentic AI as top attack vector for 2026	Dark Reading
99.4%	Of CISOs had SaaS/AI security incident in 2025	GlobeNewsWire
88%	Confirmed or suspected AI agent security incidents	AI Automation Global
82%	Discovered unknown AI agents in their environment	CSA Survey April 2026
$160B	Projected AI security market by 2029	Gartner
$400M+	Raised by NHI governance startups in 2025 alone	Aegis Intel

REGULATORY LANDSCAPE

Regulatory pressure is converging from multiple jurisdictions. The EU AI Act reaches full enforcement August 2, 2026, requiring high-risk AI compliance. The US shifted toward federal preemption of state AI laws, while Colorado and California enacted their own requirements.

Regulation	Status	Date	Key Requirements
EU AI Act	Full enforcement	Aug 2, 2026	High-risk AI rules, transparency, human oversight, stop/correct controls
HIPAA Final Rule	Targeting	May 2026	Mandatory encryption, MFA, network segmentation, 72h restoration, pen testing
Colorado AI Act	Effective	Feb 2026	Impact assessments, consumer appeal rights for AI decisions
CA AI Transparency (SB 942)	Effective	2026	AI disclosure mandates, responsible system behavior
ISO 42001	Accelerating	2026-2027	83% of Fortune 500 procurement requiring alignment by 2027
NIST AI RMF 1.1	In development	2026	Critical infrastructure AI profile, cyber AI guidance
OWASP Agentic Top 10	Published	Dec 2025	Goal misalignment, tool misuse, delegated trust, memory risks

SECURITY VENDORS AND MARKET

Company	Product	Focus	Notable
Astrix Security	AI Agent Control Plane	NHI governance	$85M raised; Cisco acquisition bid
Oasis Security	Agentic Access Mgmt	Runtime access control	$75M from Sequoia, Accel
Clutch Security	NHI platform	Centralized NHI inventory	Zero-trust enforcement
Token Security	NHI platform	Posture management	Universal NHI visibility
Lakera	Lakera Guard	AI firewall	Acquired by Check Point, 98%+ detection
Prompt Armor	Prompt security	Injection defense	Real-time prompt scanning
NVIDIA	NeMo Guardrails	LLM guardrails	Open-source, programmable safety

NHI governance startups raised $400M+ in 2025. Gartner projects $160B AI security market by 2029.

06 RECOMMENDATIONS

1. Inventory every AI identity. Every agent, service account, OAuth token, and API key in a single inventory with ownership, scope, and expiration metadata. The 21% who maintain real-time inventories have a massive advantage.

2. Enforce least privilege aggressively. 97% of NHIs have excessive privileges. Minimum permissions, scoped to specific resources, with time-bound access that expires automatically. No permanent admin tokens for POC projects.

3. Rotate credentials at machine speed. 92% fail to rotate on a 90-day cycle. For AI agents, 90 days is too long. Short-lived credentials, automatic rotation, and just-in-time provisioning should be the default.

4. Monitor for behavioral anomalies. Agents behave differently from humans and from each other. Baseline each agent's normal operating pattern and alert on deviations. An agent querying databases it has never accessed before is a signal.

5. Govern the AI supply chain. The Vercel breach started at a third-party AI tool. Every AI vendor your employees connect to extends your attack surface. OAuth grants to third-party AI tools need the same scrutiny as vendor access reviews.

6. Build accountability chains. When an AI agent takes an action, there must be a clear chain: which agent, whose authority, what permissions, what purpose. Without this, incident response is impossible.

7. Prepare for EU AI Act compliance. Full enforcement begins August 2, 2026. High-risk AI systems require technical documentation, human oversight, and stop/correct controls. Start gap assessments now.

8. Adopt ISO 42001 early. 83% of Fortune 500 procurement teams plan to require ISO 42001 alignment by 2027. Early certification becomes a competitive advantage in enterprise sales.

9. Implement agent firewalls. Deploy runtime guardrails (Lakera Guard, NeMo Guardrails) to detect prompt injection, data exfiltration, and unauthorized actions before they execute.

10. Close the budget gap. 97% expect an incident, but only 6% of security budgets address it. Reallocate security spending to match the actual threat surface.

07 METHODOLOGY + SOURCES

This report synthesizes data from publicly available research reports, industry surveys, regulatory documents, vendor disclosures, and open-source ecosystem metrics. All statistics are attributed to primary sources. Projections are directional, not definitive.

PRIMARY SOURCES

Grand View Research (2024): AI Agent Market Size and Trends

MarketsandMarkets (2024): AI Agent Market Projections

Gartner (2024-2026): Enterprise Forecasts, AI Security Platforms

McKinsey Global Survey on AI (2024): Enterprise Adoption

Saviynt / Cybersecurity Insiders (April 2026): CISO AI Risk Report

SANS Institute (April 2026): State of Identity Threats Survey

Palo Alto Unit 42 (2026): Global Incident Response Report

Cloud Security Alliance (2026): AI Agent Security Survey

OWASP (2026): Top 10 for Agentic Applications

GitHub / Microsoft: Copilot Adoption Data

Stack Overflow Developer Survey (2024)

BCG / Harvard (2023): AI Productivity Study

Deloitte (2024): State of Generative AI

Security Boulevard (2026): Agentic AI Security Report

Bessemer (2026): Securing AI Agents Analysis

CrowdStrike (2026): Global Threat Report

Wiz Research: Prompt Injection Statistics

Dark Reading: Agentic AI Attack Surface Analysis

08 ABOUT THE AUTHOR

Imiel Visser is a software engineer, writer, and AI researcher based in South Africa. He publishes original research and analysis on AI agents, AGI architecture, cybersecurity, and developer tools at imiel.dev.

Published research includes work on AGI architecture (The Subconscious Layer), developer psychology in the AI era (Meta-Imposter Syndrome), and human-computer communication frameworks (Descriptive Language Processing).

LET'S CONNECT

Questions about this report? Discuss AI agent strategy for your organization?

Book a call: cal.com/imiel

Read more: imiel.dev

Follow: x.com/imiel_visser

Connect: linkedin.com/in/imiel