



BIGID RESEARCH REPORT · 2026

The Ungoverned Enterprise:

AI Risk, Agentic Access, and the Data Security Gap in 2026

AI adoption is nearly universal. AI governance is nearly nonexistent. The gap between them is where your most sensitive data is being exposed — by shadow AI tools, autonomous agents, and machine identities that operate at scale with no human oversight and no controls.

01 The Adoption Gap

72% of enterprises have AI in production. Only 6% have updated their governance frameworks to cover it.

02 The Shadow AI Problem

67% of executives say their org already had a breach tied to unsanctioned AI. Most didn't know until it was too late.

03 The Non-Human Identity Explosion

80 machine identities for every human one — and the governance tools built for people have no answers for agents.

04 Agentic AI & the New Attack Surface

Agents chain data across systems in a single task. No single access event triggers a DLP alert. Traditional controls are blind to it.

05 AI Risk as a Business Resilience Issue

94% of security leaders say AI is the #1 driver of cyber risk change. Boards are accountable. Bain: most orgs need 2× current security spend.

06 The Data Control Plane for Responsible AI

Google, Databricks, Microsoft all racing to own the agentic control plane. The data governance layer is the one they all depend on.

07 The Path Forward

What governed AI looks like — four capabilities, a maturity model, and the \$1.9M financial case for investing now.

67%

of executives say their organization has already suffered a breach tied to unsanctioned AI tools

45%

of employees are now regular AI users on corporate devices — up 3× in twelve months

Hours

now separates vulnerability discovery from working exploit — down from weeks or months before frontier AI

Sources: Verizon 2026 DBIR · IBM Cost of a Data Breach 2025 · KPMG 2026 Cybersecurity Report · Forrester Top Threats 2026 · McKinsey Agentic Enterprise 2026 · Gartner CIO Survey 2026 · Bain & Company 2026 · EY Global Cybersecurity Leadership Insights 2026 · WEF Global Cybersecurity Outlook 2026 · Cloud Security Alliance · GitGuardian · Bessemer Venture Partners · Cyberhaven · ISACA/WEF · F5 Networks

CHAPTER 01

Speed Without Governance

Enterprise AI deployment has outpaced the ability to govern it. Most organizations can describe how they are using AI. Very few can describe what that AI can actually access.

McKinsey's Global AI Survey found **72% of enterprises have at least one AI workload in production** — up from 20% in 2020. Gartner's annual survey of nearly 1,000 CIOs found cybersecurity ranked as the number one CIO priority for the third consecutive year, with operationalizing AI in second. These two priorities are inseparable: every AI deployment expands the security perimeter that same team is trying to harden.

72%

of enterprises have AI in production — up from 20% in 2020

McKinsey Global AI Survey, 2026

97%

of executives say their company deployed AI agents in the past year

Writer / Workplace Intelligence, 2026

6%

updated governance frameworks for those agents

ISACA/WEF, 2026

AI Deployment vs. Governance Readiness

% of organizations — 2026 cross-industry surveys

DEPLOYMENT

Enterprises with AI in production **72%**

Execs who deployed AI agents in 2025 **97%**

CIOs investing in AI & ML **67%**

GOVERNANCE & READINESS

Have formal AI security policy **23%**

Updated governance for agents **6%**

Highly ready to secure AI **2%**

Sources: McKinsey Global AI Survey (2026), Writer/Workplace Intelligence (2026), Gartner CIO Survey (2026), ISACA/WEF (2026), Metomic (2026)

"Governance maturity, not budget, is the primary constraint separating organizations that contain AI-related exposure from those that do not."

ISACA / Accenture / World Economic Forum — convergent finding, 2026

• HOW BIGID HELPS

- **Discover what data AI systems can reach** — automated discovery and classification across cloud, on-prem, and hybrid environments before governance gaps form.
- **Map AI data access in real time** — BigID's data access intelligence shows exactly which AI tools, agents, and integrations touch sensitive data, and flags overpermissioned access.
- **Enforce policy from day one** — governance frameworks that scale with AI deployment rather than running months behind it.


CHAPTER 02

The Shadow AI Problem

Sensitive data is flowing into AI systems your security team cannot see. The Verizon 2026 DBIR — covering 31,000 incidents across 145 countries — documents shadow AI as a material breach contributor for the first time.

Employee AI Usage on Corporate Devices

Verizon 2026 Data Breach Investigations Report



Source: Verizon 2026 Data Breach Investigations Report

What Data Employees Send to Unauthorized AI

Top data types in DLP events targeting generative AI tools —
Verizon 2026 DBIR



Source: Verizon 2026 DBIR — Figure 65

Shadow AI became the **3rd most common non-malicious insider action** in DLP datasets — a fourfold increase year over year.

Verizon 2026 DBIR

• HOW BIGID HELPS

- **Detect what shadow AI tools can see** — BigID scans data environments to identify where sensitive data is exposed to unsanctioned AI tools and browser-based access points.
- **Classify what's at risk** — PII, PHI, source code, financial data, and regulated content automatically tagged so exposure is understood before it becomes a breach.

39%

of employee AI interactions involve sensitive data

Cyberhaven, 2026

1 in 5

breaches had shadow AI as a contributing factor

IBM, 2025

67%

of executives say their company already suffered a breach from unsanctioned AI

Writer / Workplace Intelligence, 2026

Average breach detection time is **247 days**. Shadow AI exposure typically surfaces long after the damage has compounded.


CHAPTER 03

The Non-Human Identity Explosion

The governance model broke when the ratio flipped. Most enterprises now have 80 machine identities for every human one — and the tools built to govern people have no answers for agents.

Human vs. Non-Human Identity Ratio


Average enterprise — KPMG 2026 Cybersecurity Report



Source: KPMG 2026 Cybersecurity Report

NHI Population Growth, 2023–2025

Year-over-year — Entro Labs via CSA, 2026



Source: Entro Labs H1 2025 Research via Cloud Security Alliance

The Cloud Security Alliance describes the current state as a non-human identity governance vacuum. AI agents acquire permissions dynamically, spawn sub-agents, invoke APIs, and chain actions across dozens of systems. A single compromised NHI credential has categorically greater blast radius than a compromised human account.

16%

don't track AI identity creation at all
CSA, 2026

28.6M

hardcoded secrets added to GitHub in 2025, up 34% YoY
GitGuardian, 2026

92%

say managing AI agents is the defining security skill of the next 5 years
KPMG, 2026

When an AI agent acts, most audit systems log it under the human user's credentials. **The audit trail points to a person who made no decision.**

Cloud Security Alliance NHI Governance Whitepaper, 2026

• HOW BIGID HELPS

- **Govern NHI data access** — BigID identifies what data every agent, service account, and machine credential can reach, and enforces least-privilege at the data layer.
- **Maintain traceable audit records** — every data access event attributed to the correct identity — human or non-human — for investigation, compliance, and board reporting.
- **Reduce blast radius proactively** — continuous access reviews flag overpermissioned NHIs before a compromised credential becomes a full-scale data breach.


CHAPTER 04

Agentic AI and the New Attack Surface

Traditional DLP was built to catch humans moving data. Agents don't move data the way humans do — they chain it across systems at machine speed, in combinations no single access policy was designed to catch.

Top Security Threats Identified for 2026

% identifying as top concern — multi-source synthesis



Sources: Dark Reading via Bessemer VP; Forrester Top Cybersecurity Threats 2026; WEF Global Cybersecurity Outlook 2026

An agent with access to email, a code repository, and a customer database can correlate data across all three and transmit it externally in a single task.

No single access event would have triggered a DLP alert.

CSA / Aona AI Research, 2026

Forrester's Five Agentic Threat Categories

- Near-autonomous **nation-state attacks** at machine speed
- **Rogue AI agents** entering via browser hooks and inbox access
- **AI software supply chain** exposure via open-source agentic frameworks
- **AI identity sprawl** — machine credentials multiplying without governance
- **Digital sovereignty fragmentation** creating hidden infrastructure risk

Source: Forrester Top Cybersecurity Threats 2026

● HOW BIGID HELPS

- **See what agents can chain together** — BigID maps the sensitive data footprint across all connected systems, so you understand the true blast radius before an agent is compromised.
- **Policy enforcement at the data layer** — controls applied to the data itself, not just the perimeter, so chained access events don't create undetected exfiltration paths.
- **MCP server governance** — BigID's MCP integration lets AI systems access data with full classification context, enforcing what agents are permitted to see and act on.


CHAPTER 05

AI Risk Is Now a Board-Level Business Risk

In April 2026, a single AI model changed the calculus of enterprise risk permanently. The board-level conversation about cybersecurity has not been the same since.

How AI Has Changed the Threat Landscape

% of security leaders — WEF Global Cybersecurity Outlook 2026 & EY GCLI 2026



Sources: WEF Global Cybersecurity Outlook 2026; EY Global Cybersecurity Leadership Insights Study 2026; Fortinet CISO Survey 2025

Bain: "AI does not create new vulnerabilities — it exposes existing ones." Most organizations need to increase cybersecurity spending by up to 2×. Annual 10% increases fall far short of what the threat now demands.

Bain & Company: Claude Mythos and the AI Cybersecurity Wake-Up Call, May 2026

APRIL 7, 2026 · THE MYTHOS MOMENT

Anthropic announced Claude Mythos Preview — then immediately restricted its public release. It autonomously discovers and chains zero-day exploits across enterprise systems, completing in hours what previously required elite human researchers.

Hours

exploit window, down from weeks. Defenders and attackers now operate at the same speed.

1,000s

of zero-days already identified. Project Glasswing gave AWS, Apple, Google, and Microsoft early access to patch first.

Mythos did not create a new category of risk — it proved one already present in most enterprise environments. AI systems operating autonomously with access to sensitive data, without governance controls or oversight, are the default state today. The question is whether your organization can see it.

2×

cybersecurity spend increase required to match AI-enabled threat velocity

Bain & Company, 2026

73%

personally hit by AI-enabled fraud or cyber incident in the past 12 months

WEF, 2026

• HOW BIGID HELPS

- **Close the visibility gap before attackers exploit it** — 36% of enterprise assets sit in EY's vulnerability zone. BigID eliminates the data visibility portion of that gap with continuous discovery and classification.
- **Board-ready risk reporting** — exposure translated into business impact metrics boards and CIOs can act on, not technical dashboards. Proof of control before the next board meeting.

CHAPTER 06

The Data Control Plane for Responsible AI

Google, Databricks, and Microsoft are all racing to own the agentic enterprise control plane. Every one of those architectures depends on the same foundational layer: governed, classified, trusted data. That layer is BigID.

The Control Plane Race

The most significant architectural debate in enterprise AI is who governs the layer agents run on. Bain's analysis of Google Cloud Next 2026 found Google repositioning toward **"a full agentic control plane — governed data architecture and NHI management at the center."** Databricks made the same argument: agents are only as trustworthy as the governance layer beneath them. The CSA frames the requirement: **"enforce boundaries in dynamic autonomous environments and continuously validate trust at scale."** Both point to the same foundation: classified, governed data.

THE AGENTIC ENTERPRISE STACK	
AI Agents & Applications	Claude, Copilot, Gemini, custom agents
Orchestration & MCP Layer	Tool routing, context, agent coordination
Identity & Access Layer	Human and non-human identity, permissions
Data Control Plane ← BigID	Discovery · Classification · Governance · Audit
Data Infrastructure	Cloud, on-prem, hybrid data stores

Sources: Bain/Google Cloud Next 2026; Bain/Databricks Summit 2026; CSA Agentic Control Plane 2026

Responsible AI Requires a Trust Layer

McKinsey's 2026 AI Trust Maturity Survey found an overall responsible AI maturity score of just 2.3 out of 4 — with governance and agentic controls lagging hardest. The core finding: **enterprises can build agents faster than they can build accountability for them.**

74%

of all AI economic value captured by the 20% who invest in governance infrastructure
PwC, 2026

97%

of organizations breached by AI attacks lacked proper AI access controls at the time
IBM, 2025

55%

increase in documented AI incidents in 2025 — 362 total, up from 233 in 2024
Stanford HAI, 2026

"Supporting AI no longer just means making data available. It means exposing a **governed context architecture that agents can trust.**"


Bain & Company — Google Cloud Next 2026 analysis, April 2026

- **BIGID AS THE DATA CONTROL PLANE**
- **The trust layer every agentic stack depends on** — discovery, classification, and access governance that makes every agent above it trustworthy by design. Not a wrapper. Infrastructure.
- **MCP-native governance** — BigID's MCP server gives AI agents classified, permissioned data access with full audit trails. Agents act on governed context, not raw uncontrolled data.
- **Responsible AI as infrastructure, not a checkbox** — automated policy enforcement at runtime, continuous monitoring, and board-ready proof of control built into every AI deployment.


CHAPTER 07

What Governed AI Looks Like

The organizations pulling ahead are not moving slowest. They built data governance into their AI deployment before scaling it. Governance maturity — not budget — is the differentiator.



Most enterprises in 2026 are at Stage 1. The gap to Stage 2 is smaller than it appears — and the business case is measured in millions.



Four Capabilities of the Governed Enterprise

- 1 Know what data exists and where**
Discovery and classification at scale. You cannot govern what you haven't found.
- 2 Understand what AI systems can access**
Data access intelligence mapping which agents and identities reach sensitive data.
- 3 Enforce least-privilege at the data layer**
Reduce overpermissioned access before an agent exploits it — not after.
- 4 Maintain auditable records of agent actions**
Monitoring that covers non-human identities with the same rigor as human users.

- **HOW BIGID HELPS**
- **The only platform that covers all four capabilities** — discovery, access intelligence, policy enforcement, and continuous compliance in a single AI-native data security platform.
- **Move from Stage 1 to Stage 2 fast** — BigID's agentless discovery deploys in days, not months, giving security teams immediate visibility into what AI systems can reach.

The window to act is narrowing. The organizations that govern their data now will be structurally ahead. The ones that wait are already behind.

Every week without AI data governance is a week of compounding exposure. Shadow AI is active. Non-human identities are multiplying. Agentic systems are making decisions about your most sensitive data without oversight. Frontier AI has crossed a threshold and the capability to exploit that exposure is no longer limited to nation-states.

Why the Status Quo No Longer Holds

The findings in this report converge on the same conclusion from five directions. The adoption gap is structural and will not self-correct. Shadow AI is already in your environment — 67% of executives say their organization has already experienced a breach tied to unsanctioned AI tools. Non-human identities outnumber humans 80 to 1 and most are ungoverned. Agentic systems can chain data across your most critical systems in a single task. And frontier AI has compressed the window from vulnerability to exploit from months to hours.

The IIA's 2026 Risk in Focus survey — 4,000 chief audit executives across 131 countries — ranked cybersecurity the number one global risk for the fifth consecutive year. Business resilience ranked third. They are the same problem: **an organization that cannot see what its AI systems can access cannot claim to be resilient.**

Boards need to stop asking "Are we secure?" and start asking: **"What scenarios could completely compromise us?"** That question starts with knowing what your data is, where it lives, and what can reach it.

RSAC 2026 — HackerOne CEO Kara Sprague, cited in SecurityScorecard analysis

EY found **36% of enterprise assets sit in a vulnerability zone** — below-average visibility and below-average security coverage. The data layer is the largest and least-governed portion of that gap.

EY Global Cybersecurity Leadership Insights Study, 2026

Why BigID Is Different

Most security tools were built for a world where identities are human, data is static, and access patterns are predictable. That world no longer exists. BigID was built for the one that does.

1 Data-first, not perimeter-first

BigID governs at the data layer — the only layer that matters when agents, NHIs, and shadow AI bypass the perimeter by design.

2 Built for non-human identities

BigID's agentic access governance covers service accounts, API keys, AI agents, and MCP-connected tools with the same rigor as human access reviews.

3 Discovery that moves at AI speed

Agentless deployment. Continuous classification. Real-time access intelligence. BigID closes the visibility gap in days, not months.

4 Proof of control for boards and regulators

Continuous compliance, audit-ready records of every data access event, and business-impact metrics that translate exposure into language boards use to make risk decisions.

See what your AI systems can reach.

BigID gives CISOs and CIOs the visibility, governance, and proof of control to deploy AI with confidence — and defend against it.

bigid.com

About This Report

This report synthesizes publicly available research from third-party organizations. BigID did not conduct a primary survey. All statistics are sourced and attributed. Produced by BigID's marketing team, reflecting published research as of June 2026.

Primary Sources

Verizon 2026 Data Breach Investigations Report

31,000+ incidents and 22,000 confirmed breaches across 145 countries. Covers November 2024 through October 2025.

KPMG 2026 Cybersecurity Report

Interviews with 20+ KPMG cyber leaders and senior executives. Identifies non-human identities as one of eight critical risk priorities for 2026.

McKinsey: Securing the Agentic Enterprise, 2026

Survey of cybersecurity decision-makers in financial services and software. Published March 2026.

EY Global Cybersecurity Leadership Insights Study, 2026

Survey of 800+ cybersecurity leaders analyzing 475 asset types. Source of the 36% "vulnerability zone" finding.

Gartner C-level Communities CIO Survey, 2026

Annual survey of nearly 1,000 CIOs. Third consecutive year cybersecurity ranked as the number one CIO priority.

GitGuardian State of Secrets Sprawl, 2026

Annual repository secrets exposure analysis. Documents 28.65 million hardcoded secrets added to GitHub in 2025 — a 34% YoY increase.

McKinsey Global AI Survey, 2026 · Writer / Workplace Intelligence: Enterprise AI Adoption Survey, 2026 (n=2,400) · Cyberhaven Research, 2026 · ISACA/Accenture/WEF convergent findings, 2026 · Anthropic: Assessing Claude Mythos Preview's Cybersecurity Capabilities, April 2026

IBM Cost of a Data Breach Report, 2025

Annual study of financial impact of data breaches. Includes dedicated analysis of AI-related breach costs and shadow AI.

Forrester Top Cybersecurity Threats 2026

Annual forward-looking threat analysis. The 2026 edition identifies five categories of emerging risk driven by agentic AI deployment.

Bain & Company: Claude Mythos and the AI Cybersecurity Wake-Up Call, 2026

Analysis of Claude Mythos Preview's cybersecurity implications and the required enterprise response.

WEF Global Cybersecurity Outlook 2026

Fifth annual WEF cybersecurity survey in collaboration with Accenture. 94% and 87% AI risk findings sourced here.

Cloud Security Alliance: NHI Governance Whitepaper, 2026

Analysis of the non-human identity governance gap. Draws on CSA token sprawl analysis and Entro Labs H1 2025 research.

Bessemer Venture Partners: Securing AI Agents, 2026

Framework for CISO response to agentic AI risk. Published March 2026. Includes Dark Reading poll on top security threats.

About BigID

BigID is the leading platform for data security, privacy, compliance, and AI governance. BigID helps organizations discover, classify, and control their most sensitive data across cloud and on-premises environments — enabling security teams to see and govern what AI systems can access before exposure becomes a breach. bigid.com