



Mythos Changed the Attack Surface. BigID Secures What's Behind It.

EXECUTIVE BRIEF — AI RISK & BUSINESS RESILIENCE

Anthropic's Claude Mythos found more than 10,000 critical vulnerabilities in the world's most important software before most organizations had updated their threat models. The capability is expanding. Pricing is falling. And when Mythos-class AI reaches adversaries, the decisive factor will not be perimeter defenses. It will be whether your organization already knows what data it holds, what it means to the business, and whether the right controls are in place before exposure becomes a breach, a headline, or a regulatory event.

WHAT MYTHOS ACTUALLY DEMONSTRATED

A capability shift that made legacy security assumptions obsolete

Claude Mythos Preview launched in April 2026 with capabilities that had no real precedent: it autonomously found thousands of zero-day vulnerabilities across every major operating system and browser, compressed exploit discovery from weeks to hours, and produced working attacks without human direction. In its first weeks of restricted deployment through Project Glasswing (a coalition including AWS, Apple, Cisco, CrowdStrike, Google, JPMorganChase, Microsoft, NVIDIA, and Palo Alto Networks), Mythos identified more than 10,000 high- or critical-severity vulnerabilities in the world's most systemically important software.

Anthropic withheld broad release deliberately, citing the risk of arming attackers before defenders were ready. That calculus is already shifting. In June 2026, Anthropic released Claude Fable 5 publicly and Claude Mythos 5 to an expanding group of vetted partners. Pricing has dropped significantly. Access is broadening. Anthropic has confirmed that Mythos-class capabilities will reach general availability. The question is not whether adversaries will have access to tools like this. It is when.

THE BUSINESS RESILIENCE PROBLEM

What breaks when AI-powered attackers reach your data estate

Most enterprise resilience programs were built around a containable breach scenario: an attacker gets in, moves laterally, extracts something specific, gets detected. The response playbook, the cyber insurance policy, the incident response retainer, and the board reporting structure were all designed for that model.

Mythos-class AI does not operate that way. It reasons through environments systematically, at machine speed, without fatigue. It does not need a map of your data estate. It will build one. And when it does, the exposure is not targeted. It is total: customer PII, intellectual property, source code, financial records, regulated data, model training sets, and business-critical knowledge repositories, all surfaced and indexed before a human analyst has logged the first alert.

The downstream consequences extend well beyond security operations. A breach at this scale triggers regulatory notification obligations across GDPR, CCPA, HIPAA, and sector-specific frameworks. It creates board-level liability exposure. It activates supply chain and third-party risk provisions. It generates reputational damage that outlasts any technical remediation. And it happens faster than any incident response plan written for human-speed attackers can address.

The organizations that are genuinely resilient to this threat model are not the ones with the best detection tools. They are the ones that already know what data they hold, where it lives, what it is worth to an attacker, and what the business impact of its exposure would be. That foundation has to exist before the incident, not be assembled during it.

HOW BIGID BUILDS THAT FOUNDATION

From crown jewel discovery to AI risk management to enforced controls

Crown Jewel Discovery

Resilience starts with knowing what matters most. BigID's discovery and classification engine operates at a depth and scale that generic security tools cannot match: it understands data in context, correlating sensitivity, identity, location, and business value across 100+ cloud, on-prem, SaaS, and unstructured sources. Organizations get a complete, continuously updated inventory of their data estate, including PII, credentials, source code, regulated records, and shadow data accumulating in forgotten systems, giving security and compliance teams a clear picture of which assets carry the highest business impact if compromised.

AI Risk Management

Most organizations are deploying AI faster than they are governing it. BigID maps sensitive data exposure across every AI pipeline, RAG architecture, copilot, and autonomous agent in the environment, identifying which models can reach regulated information, flagging ungoverned inputs before they enter model context, and tracking third-party AI tool risk alongside internal deployments. Where other solutions treat AI as a separate problem, BigID integrates AI risk directly into the enterprise resilience program so every new AI surface is governed from the start.

DSPM & AISPM

BigID delivers data security posture management and AI security posture management in a single platform, continuously monitoring how sensitive data is stored, accessed, and exposed across the full environment, including AI systems, agents, and pipelines. Security and compliance teams get risk-based prioritization of critical assets, automated remediation workflows, and audit-ready reporting that reflects current posture. When regulators or the board ask what was exposed and how it was governed, the answer is already documented.

Intelligent Security Controls

BigID enforces controls that go beyond who is accessing data to evaluate why, applying policy based on sensitivity, business context, and risk in real time. Unlike point solutions that address a single control layer, BigID operates across the full data lifecycle:

- Data access governance across users and AI systems
- Agentic access controls for autonomous AI workflows, evaluated against intended use and permission scope
- Context-aware DLP that accounts for sensitivity, data type, and business purpose
- Automated retention and deletion to continuously reduce attack surface
- Continuous exposure monitoring across every environment where sensitive data lives

KEY AI SECURITY USE CASES BIGID ADDRESSES

Covering the full scope of AI risk across the enterprise

Shadow AI

- Discover unsanctioned models and find unauthorized model use of sensitive data
- Find unauthorized users accessing models and vendors using unsanctioned AI
- Identify MCP-exposed AI data across the environment

AI Security & Governance

- AI security posture management and AI risk assessment
- Agentic access and security controls
- AI regulatory and control compliance including EU AI Act and AI NIST Framework

AI Prompt Security

- Monitor data sharing with Copilots and Agents
- Block sensitive data sharing by users and build and enforce guardrails
- Preserve access privileges for employees interacting with copilots and agents
- Data labeling for Copilot

Secure AI Data Prep

- Advanced data discovery and classification by risk and business context
- Searchable catalog of gen AI data for curation of training data
- Cleansing and compliance of curated data and audit of what data was used for what AI

Agent Access Security

- Agent identification, discovery, and access permission governance
 - Agent access activity monitoring and data tracking and traceability
-

WHY BIGID

Purpose-built for the complexity of the modern data estate and the AI era

Most data security tools were designed for a simpler environment: structured databases, known data types, predictable access patterns. The modern enterprise data estate is none of those things. It spans hundreds of sources, accumulates data faster than any team can manually govern, and now includes AI systems that access and surface sensitive information in ways no legacy tool was built to track.

BigID was purpose-built for this complexity. Where others scan metadata, BigID reads data directly, using ML-based classification to understand what data actually is and what it means in business context, not just what it is labeled. Where others deliver point solutions for DSPM or DLP or privacy, BigID delivers a unified platform that connects discovery, classification, risk, governance, and controls in a single system of record for the data estate.

That unified foundation is what makes BigID different in a Mythos-era threat environment. When an AI-powered attacker can find, correlate, and exploit sensitive data at machine speed, the only meaningful defense is a platform that already has that picture, with the controls to act on it. BigID is the only platform named a Forrester Wave Leader in Sensitive Data Discovery and Classification that also covers AI risk, DSPM, AISPM, access governance, DLP, and privacy automation in one integrated system.

Resilience in the age of AI-powered attacks is not built at the perimeter. It is built in the data estate, before the breach, before the regulator calls, and before the board asks what was exposed. BigID gives organizations the visibility, governance, and controls to get there first.

CUSTOMER VALIDATION

Proof in Action

"BigID gives me better visibility into sensitive data, helps prioritize security protections, reduces attack surfaces, strengthens compliance, and increases operational efficiency overall — a strategic pillar of our AI-First Cybersecurity Transformation"

- CISO, Global Healthcare Company

"We needed to automate processes and data management across systems for the data of a few million customers across a lot of systems and data. BigID was the one solution that did this in the most efficient and sophisticated way - and had more use cases we could add on moving forward."

- Chief Privacy Officer, Global Telecoms Company

ANALYST RECOGNITION

Validated by Industry Leaders

- 2025 Company of the Year for AI Governance - Frost & Sullivan
- Named a Leader across all DSPM research including Omdia DSPM Universe, CB Insights for DSPM, TAG, GigaOm, Frost & Sullivan
- Named a Leader in all Privacy Management evals including Forrester Privacy Management Wave & IDC Privacy Compliance MarketScape
- BigID named a Leader across 4 GigaOm evals: DSPM, Data Security Platforms, Unstructured Data Management, Data Access Governance
- Named in Gartner's 2025 DLP Market Guide
- Named in Gartner's 2025 AI TRiSM Market Guide
- Represented 30+ times in Gartner's 2025 Hype Cycles for Security, Privacy, AISPM

FREE RISK ASSESSMENT

Discover Your Data & AI Risk

An agentless, cloud-native assessment to enable security, privacy, and compliance: pinpoint where sensitive data and AI risks live - and how to fix them.

What You Get:

- Complementary 2 week engagement on real data to gain visibility on risk across data and AI
- Executive-ready report identifying high-risk data and AI across cloud environments with next-step recommendations

NEXT STEPS

Ready to Connect the Dots in Data & AI?

Schedule your free risk assessment or ask for a deeper demo across DSPM, DLP, AI SPM, Privacy Automation, AI Risk, and Data Lifecycle Management

► www.bigid.com/demo

► www.bigid.com/risk-assessment

► info@bigid.com



BigID Next

"The features they offer are the moat that they have built around the product. This is what makes them the market leader in Data Privacy and Protection."



IT Security & Risk Management Associate - Banking

Gartner
Peer Insights.

Gartner® and Peer Insights™ are trademarks of Gartner, Inc. and/or its affiliates. All rights reserved. Gartner Peer Insights content consists of the opinions of individual end users based on their own experiences, and should not be construed as statements of fact, nor do they represent the views of Gartner or its affiliates. Gartner does not endorse any vendor, product or service depicted in this content nor makes any warranties, expressed or implied, with respect to this content, about its accuracy or completeness, including any warranties of merchantability or fitness for a particular purpose.



BigID Next

"Exceptionally powerful and easy to use tool that integrated smoothly into our environment and delivered value from day one. The BigID team is fabulous to work with."



Senior Vice President & Chief Data Officer - Banking

Gartner
Peer Insights.

Gartner® and Peer Insights™ are trademarks of Gartner, Inc. and/or its affiliates. All rights reserved. Gartner Peer Insights content consists of the opinions of individual end users based on their own experiences, and should not be construed as statements of fact, nor do they represent the views of Gartner or its affiliates. Gartner does not endorse any vendor, product or service depicted in this content nor makes any warranties, expressed or implied, with respect to this content, about its accuracy or completeness, including any warranties of merchantability or fitness for a particular purpose.