AN ARXNIMBUS POSITION PAPER

THE SKILL FLOOR JUST DROPPED TO ZERO

What agentic ransomware means for
how cyber risk gets priced

31 SEC

TO REWRITE AND REDEPLOY THE EXPLOIT.

The AI-driven recovery loop Sysdig documented in the first known case of fully agentic ransomware — run end to end, with no human at the keyboard.

READ THE PAPER

ArxNimbus
Secure the business.

#ProtectYourEBITDA — it's not magic. It's math.

ArxNimbus.com

AN ARXNIMBUS POSITION PAPER · ENTERPRISE / BOARD EDITION

The Skill Floor Just Dropped to Zero:
What Agentic Ransomware Means for Exposure Pricing

By R. David Moon, CEO & Co-Founder, and Andrew Patterson, COO · ArxNimbus

The thesis in one line. Agentic AI has removed the last variable that made attacker skill a defensible input to risk pricing. What's left to price is exposure itself — in dollars, not adversary tiers.

1 Agentic ransomware is a step change. Most risk models haven't caught up.

For as long as insurers and CISOs have modeled cyber risk, one variable has quietly done a lot of work: attacker skill. Sophisticated actors reach deeper into an environment; less capable ones stall out, get caught, or simply can't execute the harder steps of an intrusion. Risk models, premiums, and board-level "how bad could this get" conversations have leaned on that gradient for years.

Skill was scarce. Scarcity was something you could price. That is the assumption threat-actor tiering and most exposure models were built on — and it just took a direct hit, in production, against a real victim, not in a lab.

On July 6, 2026, CyberScoop reported that Sysdig researchers had documented what they believe is the first fully agentic ransomware operation. That claim, and the mechanics behind it, are worth taking seriously.

600+ distinct, purposeful payloads run in rapid succession.
31 seconds[1] from a failed exploit to a working one.
And a threat actor with no prior track record ran the entire kill chain end to end.

2 The case, and why the details matter more than the headline.

The attack, run in late June 2026 by a financially motivated actor Sysdig tracks as JadePuffer, spanned reconnaissance, credential theft, lateral movement, persistence, encryption, destruction, and delivery of the ransom note itself. According to Sysdig's Michael Clark, senior director of threat research, the operation was "driven end-to-end by the model's own decision-making, rather than a human at the keyboard."

The agent exploited a known Langflow vulnerability, CVE-2025-3248, to gain initial access to a production server running MySQL and Alibaba Nacos. When a Nacos backdoor deployment failed, the agent read the error, switched its method from subprocess calls to direct library imports, and redeployed a working payload 31 seconds later — a recovery loop Clark said moved "at a speed no human matches." Sysdig also found the agent had drawn on API keys for

[1] Matt Kapko, "Sysdig clocks first documented case of agentic ransomware," CyberScoop, July 6, 2026, reporting Sysdig threat research and comments from Michael Clark, senior director of threat research.

OpenAI, Anthropic, DeepSeek, and Gemini, evidence of multiple models at work across the operation.

A person still mattered. Clark was clear on this: someone provisioned the command-and-control server and the staging server for stolen data, chose the victim, and supplied root MySQL credentials obtained through a separate, prior compromise. This wasn't an autonomous attacker picking its own targets from scratch. It was a human operator who no longer needed the deep technical bench a ransomware operation used to require.

Clark put the implication plainly: "The skill floor for running a full ransomware operation just dropped to whatever it costs to run an agent."

3 Where sophistication-based pricing stalls: activity without a number.

Threat-actor profiling improves prioritization and helps analysts pattern-match. But if it stops there — if it never translates into financial terms — it leaves leadership with a more sophisticated version of the same incomplete answer. The gap is precise:

What threat-actor profiling tells you	What it leaves unpriced
✓ how experienced the actor appears	✕ how much exposure it creates, in dollars
✓ which techniques and tools it favors	✕ how much investment removes it
✓ how advanced the campaign looks	

The board does not need to become fluent in adversary taxonomy. The CFO does not need another threat-actor tier. They need a defensible way to understand business exposure, in the same units they use for every other decision they make.

> ***"Sophisticated" was never a number the board could act on. It was a proxy — and the proxy just broke.***

That sentence is the whole problem. A threat-actor tier is easy to report. Dollar exposure is hard to prove. And without a defensible number, every cyber budget conversation defaults to faith, fear, or last year's number plus ten percent.

And the exposure that goes unpriced is not marginal. ArxNimbus modeling across our 4,000+ company risk-profile library shows that enterprises routinely carry unquantified cyber exposure exceeding 25% of EBITDA — exposure no threat-actor tier is built to surface.[2]

[2]ArxNimbus data foundation: 4,000+ company risk-profile library, 60,000+ cyber-incident loss database, 47,000+ threat-vulnerability pairings, 32 primary authoritative sources, 580+ industries; engine validated against documented loss events to within 7%.

4 The missing layer: financial exposure intelligence.

Financial exposure intelligence is what connects cyber, software, AI, and control data to business and financial outcomes. It answers the questions leadership actually asks: How much exposure do we carry today? Where is it concentrated? Which exposures create the greatest potential financial impact? Which investments reduce exposure most efficiently? How does exposure change after we remediate? What does it mean for EBITDA, ROI, resilience, and how we think about risk transfer?

The shift is simple to state and hard to deliver:

- **Visibility** tells you what exists.
- **Quantification** tells you what it means.
- **Financial exposure intelligence** tells you what to do next.

Done right, it produces two defensible numbers a board can govern against — built on actuarial models rather than heat maps and professional judgment:

Current financial exposure	Projected exposure after controls
what we carry today, in dollars	what changes when we invest

That is the difference between threat intelligence and decision intelligence. A threat report shows that an actor is capable. Exposure quantification shows what that capability costs your business and what action removes it — regardless of whether the actor behind the keyboard is a person, a script, or an agent.

And the payoff is measurable, not theoretical. Across the ArxNimbus client base, exposure reduction typically translates into a 1–3% average annual EBITDA improvement — on an engine validated against documented loss events to within 7%.

5 Why now — and why this gets bigger, not smaller.

Cyber risk no longer lives inside the security function. It intersects directly with operational resilience, AI governance, software supply chain risk, insurance, M&A, compliance, customer trust, and enterprise value. At the same time, executives are being asked to make faster decisions under more uncertainty — approving AI initiatives, defending budgets, reviewing coverage, answering regulators, and explaining posture to investors.

The old model of cyber reporting does not provide enough decision support for that environment. Leaders need to move from alerts, to exposure, to financial impact, to prioritized action, to *measurable reduction.*

And here is the part most organizations have not internalized yet: the same blind spot is now forming around AI. JadePuffer's own agent pulled credentials for OpenAI, Anthropic, DeepSeek, and Gemini in the course of the attack — a reminder that the AI supply chain enterprises now depend on is exactly as unquantified as the exposure the agent found waiting for it in production. The fragmented-signals-no-financial-number pattern that produced a decade

of cyber dashboards is repeating — faster — with AI systems being adopted ahead of the governance to measure them.

> **88%** of organizations now use AI in at least one part of the business.
> **8%** have a comprehensive framework to govern it.[3]
> **That 10× gap is the next exposure crisis — forming in plain sight, in exactly the pattern this paper just described.**

It is not hypothetical. Two-thirds of executives already believe their organization has suffered a data leak through unapproved AI tools.[4] The organizations that learn to quantify exposure now, in dollars, will be the ones able to govern AI risk when it matters. The ones still arguing over severity scores will not.

6 What leaders should actually demand from exposure management.

As agentic attacks become routine, the temptation is to respond with another threat-intelligence feed or another severity score. The better test is not *"Do we know how sophisticated our adversaries are?"* It is *"Can our exposure program answer the questions leadership actually needs answered?"*

- Can we quantify exposure in financial terms, independent of who or what is attacking us?
- Can we compare exposure across cyber, software, and AI risk on one scale?
- Can we prioritize remediation by business impact, not just severity?
- Can we model the projected effect of controls *before* we spend?
- Can we track exposure reduction *after* we invest?
- Can we explain all of it in terms the board and the CFO already use?

If the answer is no, the organization has more threat intelligence — but not more clarity. Boards do not need a briefing on how capable the next attacker's agent might be. They need exposure clarity. The adversary's resumé matters less than the answer.

7 Where ArxNimbus fits.

JadePuffer's agent didn't just attack infrastructure — it ran on infrastructure: API keys, model dependencies, and AI tooling most enterprises can't currently inventory, let alone price. That is the gap **Thrivaca™AIQ** was built to close: the only platform that translates AI exposure into defensible financial impact — connecting AI models, business processes, and revenue to quantified EBITDA risk. AIQ combines actuarial-grade financial exposure modeling with AI Bill of Materials (AIBOM) intelligence, so the models an agent touches, the dependencies it pulls, and the processes they support are inventoried and priced — aligned to the NIST AI RMF and ISO 42001, not another governance checklist or scoring model. Built on the same Thrivaca actuarial engine validated across insurance and enterprise markets, it delivers a single actuarial source

[3] AI adoption data per Aon; AI-governance-framework data per Economist Impact.
[4] Writer — State of Enterprise AI 2026.

of truth for AI exposure, in dollars, whether the operator on the other end is a person or an agent.

About ArxNimbus

ArxNimbus was built on a simple idea: AI and actuarial modeling, together, are the most reliable way to measure exposure in financial terms. Since 2016, we have used AI to normalize fragmented security data and actuarial science to quantify its financial impact — producing a continuous, defensible system of record for enterprise exposure.

Built in collaboration with leaders across national security, academia, and industry — including MITRE, the University of Chicago, and the U.S. Department of Defense — the **Thrivaca**™ engine reflects how risk is actually modeled and validated in the real world. **ThrivacaAIQ** applies that engine to AI exposure — bringing insurer-grade financial rigor to AI-driven decision-making.

About the authors

R. David Moon — Chief Executive Officer & Co-Founder

R. David Moon is a quant-focused cybersecurity risk specialist whose career spans Fortune 500 CIO and CISO roles, a Big Four consulting partnership, and senior software-industry leadership. A CISSP and four-year U.S. Air Force veteran with Secret clearance, he brings a 20-plus-year portfolio of high-value work across information security, privacy, IT risk management, asset management, and infrastructure — advising some of the world's most respected organizations across the United States, Latin America, and Europe. He examined the measurable value of technology capability in risk mitigation in his 2011 book *Webify*, and has served on the board of a public investment trust and on the senior executive committee of a $5B Nasdaq company.

Andrew Patterson — Chief Operating Officer

Andrew Patterson leads the delivery and optimization of ArxNimbus's cyber and AI risk-management solutions, drawing on a career that bridges public service and the private sector. In government, he served as Deputy Chief of Staff and Global Risk Advisor at the U.S. Department of Energy and as liaison to the National Security Council, shaping risk-management strategy at the highest levels of decision-making. In the private sector, he worked with PwC's Research and Insights team analyzing emerging global risks, and with leading hedge funds and commodity traders navigating complex operational and portfolio risk. Across decades of practice, he has pioneered approaches to emerging-threat analysis, cyber strategy, and enterprise risk frameworks that marry the tactical with the strategic.