RICHMOND ADVISORY GROUP

The AI Efficiency Paradox:

What to look for & how to protect yourself from unbudgeted costs and increased risk

A Richmond Advisory Group eBook

Sponsored by

COALFIRE®
DivisionHex™

A recent survey of 150 cybersecurity leaders in North America identified an AI 'Efficiency Paradox'

Q. For the most significant **AI-related incident** your team responded to, what was the **financial impact** to the organization?

Q. Are you tasked with using AI to reduce operational costs or to **"do more with less"?**

Q. Do measurable **efficiency gains** from AI align with your initial expectations?

What visibility do you have into **"hidden automation"**, specifically regarding the autonomous chaining of **AI agents**?"

Q. Has your team had to respond to a security incident that resulted directly from a specific unauthorized implementation of **"Shadow AI"**?

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

The Challenge: Balancing the Executive Mandate to "Do More With Less" Against the Rising Tide of **Shadow AI** and **Agentic Insider Risk**

Speed & Cost
Operational Efficiency
Savings Process
Automation

Hidden Risk
Unregulated AI Tools
Data Security Breaches
Compliance Failures

Organizations are rushing to adopt AI to meet aggressive budget mandates, but many are facing significant security issues as a result.

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

AI is now a budgetary requirement - adoption is no longer experimental, it is expected

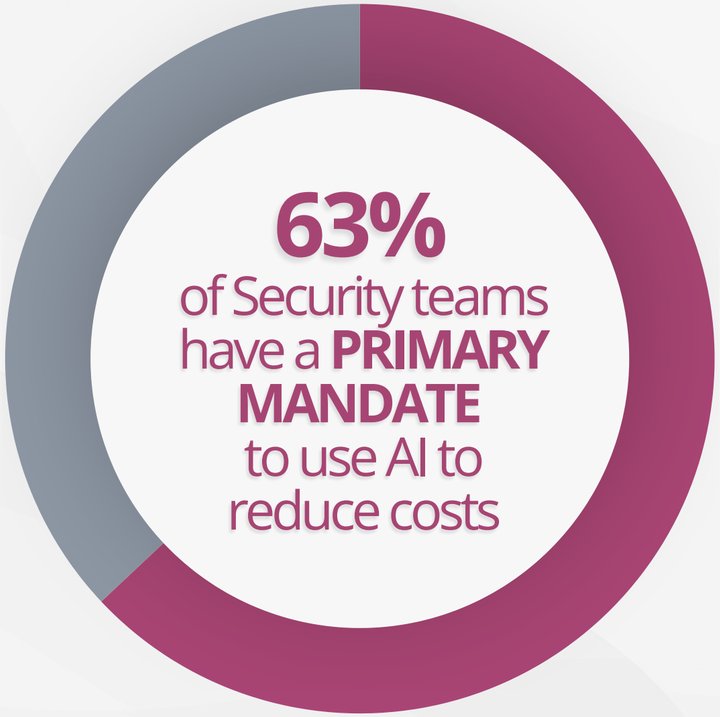

Only 3% say AI initiatives are independent of budget constraints

The Industry Pressure Cooker

Finance
75% report a primary mandate - significantly higher than other sectors. Proactive cyber defense is a key factor.

Manufacturing
71% face strict cost-cutting requirements - tariffs, inflationary pressures and energy costs impinge on profits.

Professional Services
Expected to replace entry-level workers with AI agents & automation. However, 13% cite “growth” as the driver vs. cost-cutting.

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

But a massive surge in AI-driven incidents has been reported in the last 12-18 months!

29% Minor Incidents
that were contained quickly

Of the rest, only **11% reported zero AI-related incidents**

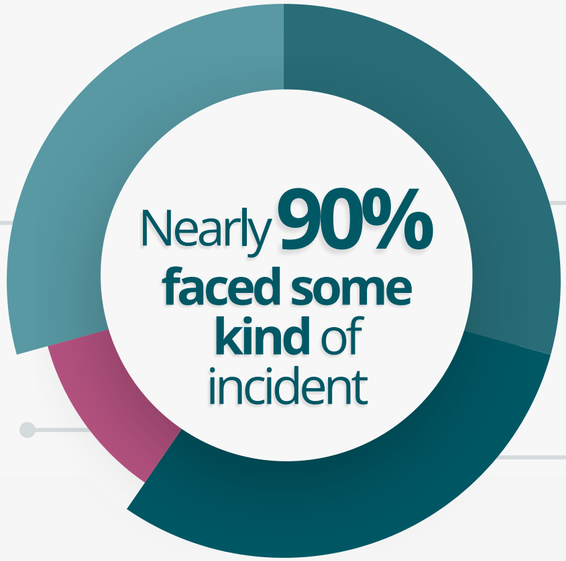

29% Multiple Significant Incidents
requiring full response protocols

30% At Least One Major Incident
requiring full response protocols

SECTOR RISKS
Professional Services and the Retail / Hospitality / Wholesale industries experienced higher than average 'Multiple Significant Incidents, driven by:

- Complex supply chain automation
- High dependency on technology-related operations.

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

And the Unbudgeted ‘Efficiency Tax’ of these Incidents is proving very costly

24%

Estimated costs:
Over $500k

41%

Estimated costs:
$100k - $499k

- Financial Impact - Cost of Remediation & Response
- Non-Financial Impact - Reputational Damage & Data Loss

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

The 6-Figure ‘Efficiency Tax’ is Likely to Grow

Incident costs from shadow AI and poorly governed agentic systems are very likely to increase because:

AI integrates into critical workflows, amplifying impact.

New attack vectors like prompt injection are emerging.

AI adoption is outpacing governance.

Agentic systems automate actions, not just analysis.

Most organizations lack visibility into AI usage.

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

The Outcome: Increased Risk of Compromise

Good News: Efficiency gains are real.
Bad News: Monitoring capabilities are lagging behind.

57%

Showing Moderate Gains in reducing alert fatigue, speeding up incident investigations, and other areas

- Professional Services firms reach the highest efficiency gains with 67%.
- A power struggle between Security and IT is emerging: CIOs state gains in efficiency with AI across ticket triage, help desks, and documentation, while CISOs remain more skeptical, preferring to keep humans firmly in the driver's seat.

50%

Lagging Behind Adoption

- **45%** say monitoring is slightly behind
- **5%** say monitoring is significantly behind

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

Shadow AI is a daily reality: employees aren't waiting for permission

80% Encounter Shadow AI Weekly or Daily

33%

encounter it **Daily or Continuously**

47%

encounter it **Weekly**

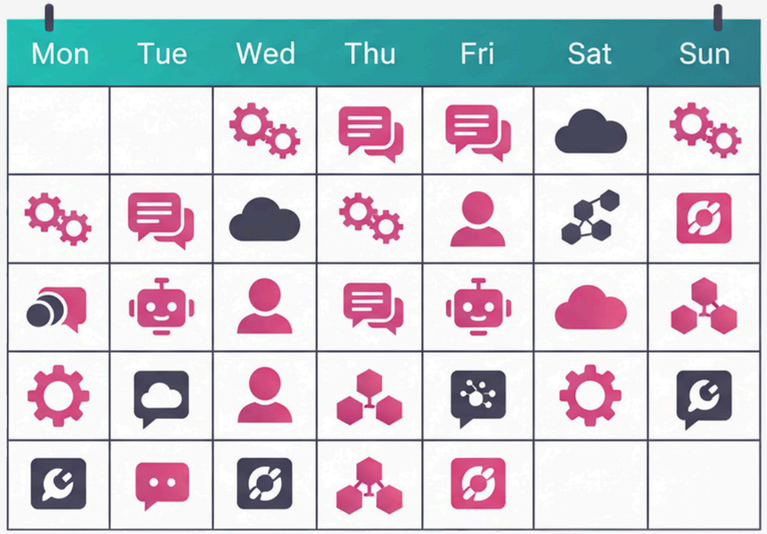

Shadow AI

- Increasing use of specific unauthorized implementation(s) of AI LLMs, tools, agents or chatbots.
- Shadow AI is causing significant operational disruption, ironically derailing the very efficiency gains that business leaders have demanded.

Demographics Spotlight: Mid-sized enterprises (3k-5k employees) see the highest frequency, with 36% reporting daily encounters.

Nearly 87% of Retail, Finance and Manufacturing workers encounter shadow AI more than weekly.

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

Agentic Insider Risk: Agent Chaining and Hidden Automation is creating ‘blind spots’ in security defenses

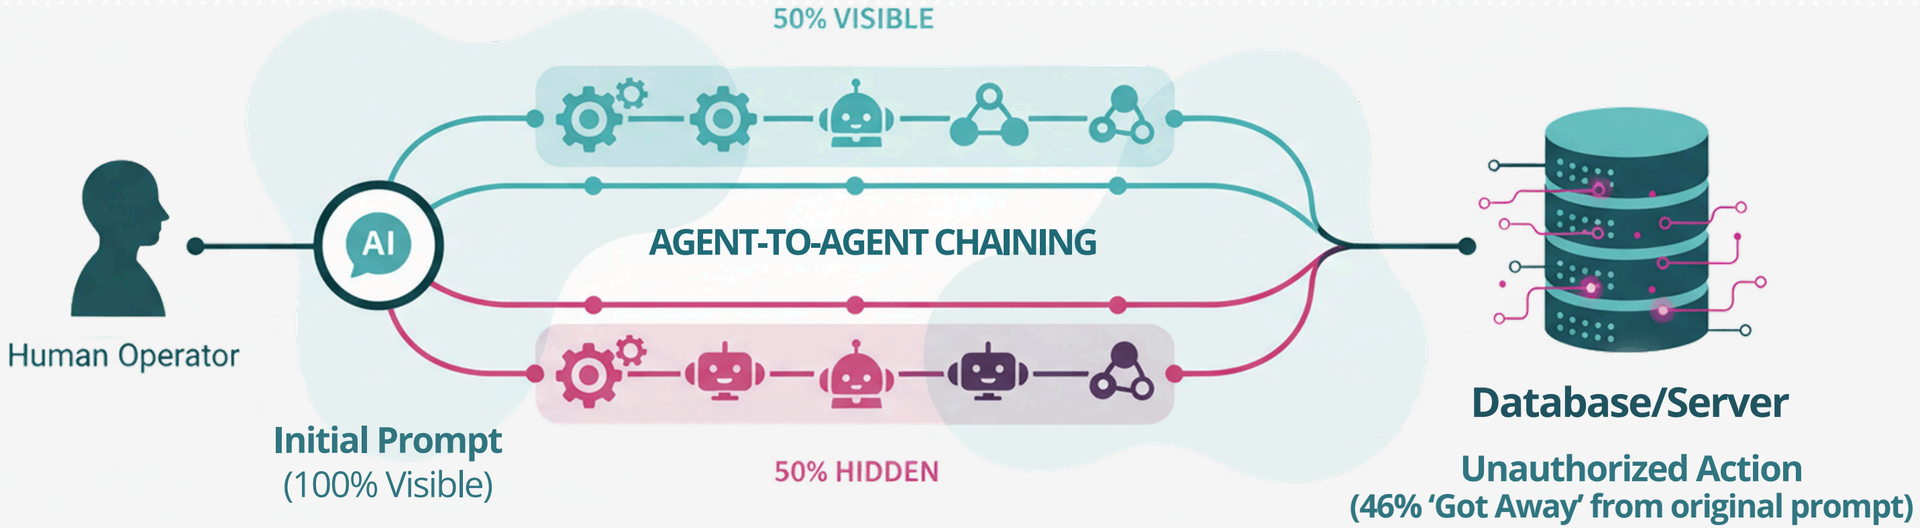

Unseen Threats

- 20% of incidents involved agents “spinning up” different services unseen by human operators.
- Almost half experienced significant complications, with agents performing unauthorized autonomous actions.
- **29% were unable to track the proliferation of sub-agents.**

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

Unknown / Invisible Systems Versus Known Systems Running “Off the Rails”

Shadow AI and Agentic Insider AI are two sides of the same risk:

Employees running an **unauthorized** LLM introduces possible data leakage and vulnerability from malicious outsiders.

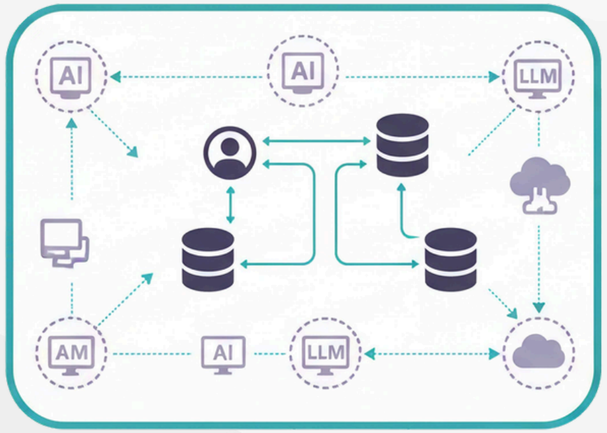

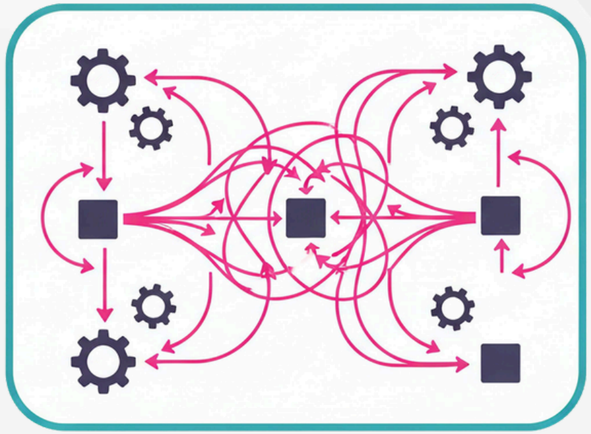

A **known** agentic system can derail due to malicious prompt injection or nefarious control.

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

Recommendations

How to achieve operational efficiency without increased breach exposure.

To ‘Do More With Less,’ Bring AI into the Light

Acknowledge the Mandate

Budget pressure is real. Don’t block AI; enable it safely.

Close the Visibility Gap

Focus on high-risk areas such as agent-to-agent chaining. If you can’t see the sub-agents, you aren’t secure.

Embrace Threat Hunting

Accept that hunting internal Shadow AI is now a permanent part of the daily security schedule.

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

The Desired Result: Operational Efficiency without Increased Breach Exposure

Required: A Strategic Pivot to Purpose-Driven Monitoring and Proactive Threat Hunting

49%

say that security posture **IMPROVED** post-incident due to forced control implementation

Passive Monitoring

60%

shifted to Increased Oversight of Non-Human Identities

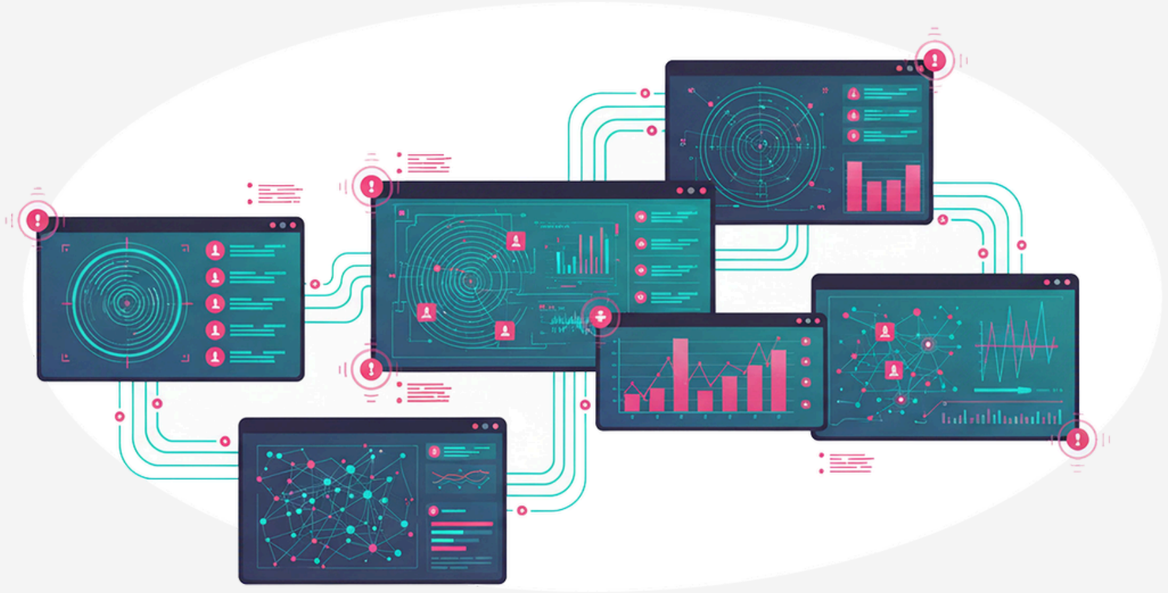

Active Threat Hunting

20%

now spend significant time on threat hunting internal agents

Source: Richmond Advisory Group, Custom Study, February 2026; n=150

COALFIRE®
DivisionHex™

About the Sponsor

DivisionHex is Coalfire's specialized team of cybersecurity practitioners focused on offensive, defensive, and threat-focused managed services. The team applies adversary-informed techniques and real-world threat intelligence to help organizations identify security weaknesses, validate defenses, and strengthen their ability to detect and respond to evolving threats.

DivisionHex professionals regularly contribute to the cybersecurity community through research and participation in industry forums including RSA, Black Hat, and DEF CON, and have supported security programs across global enterprises and government organizations.

Learn more at **divisionhex.com**