FORTRA

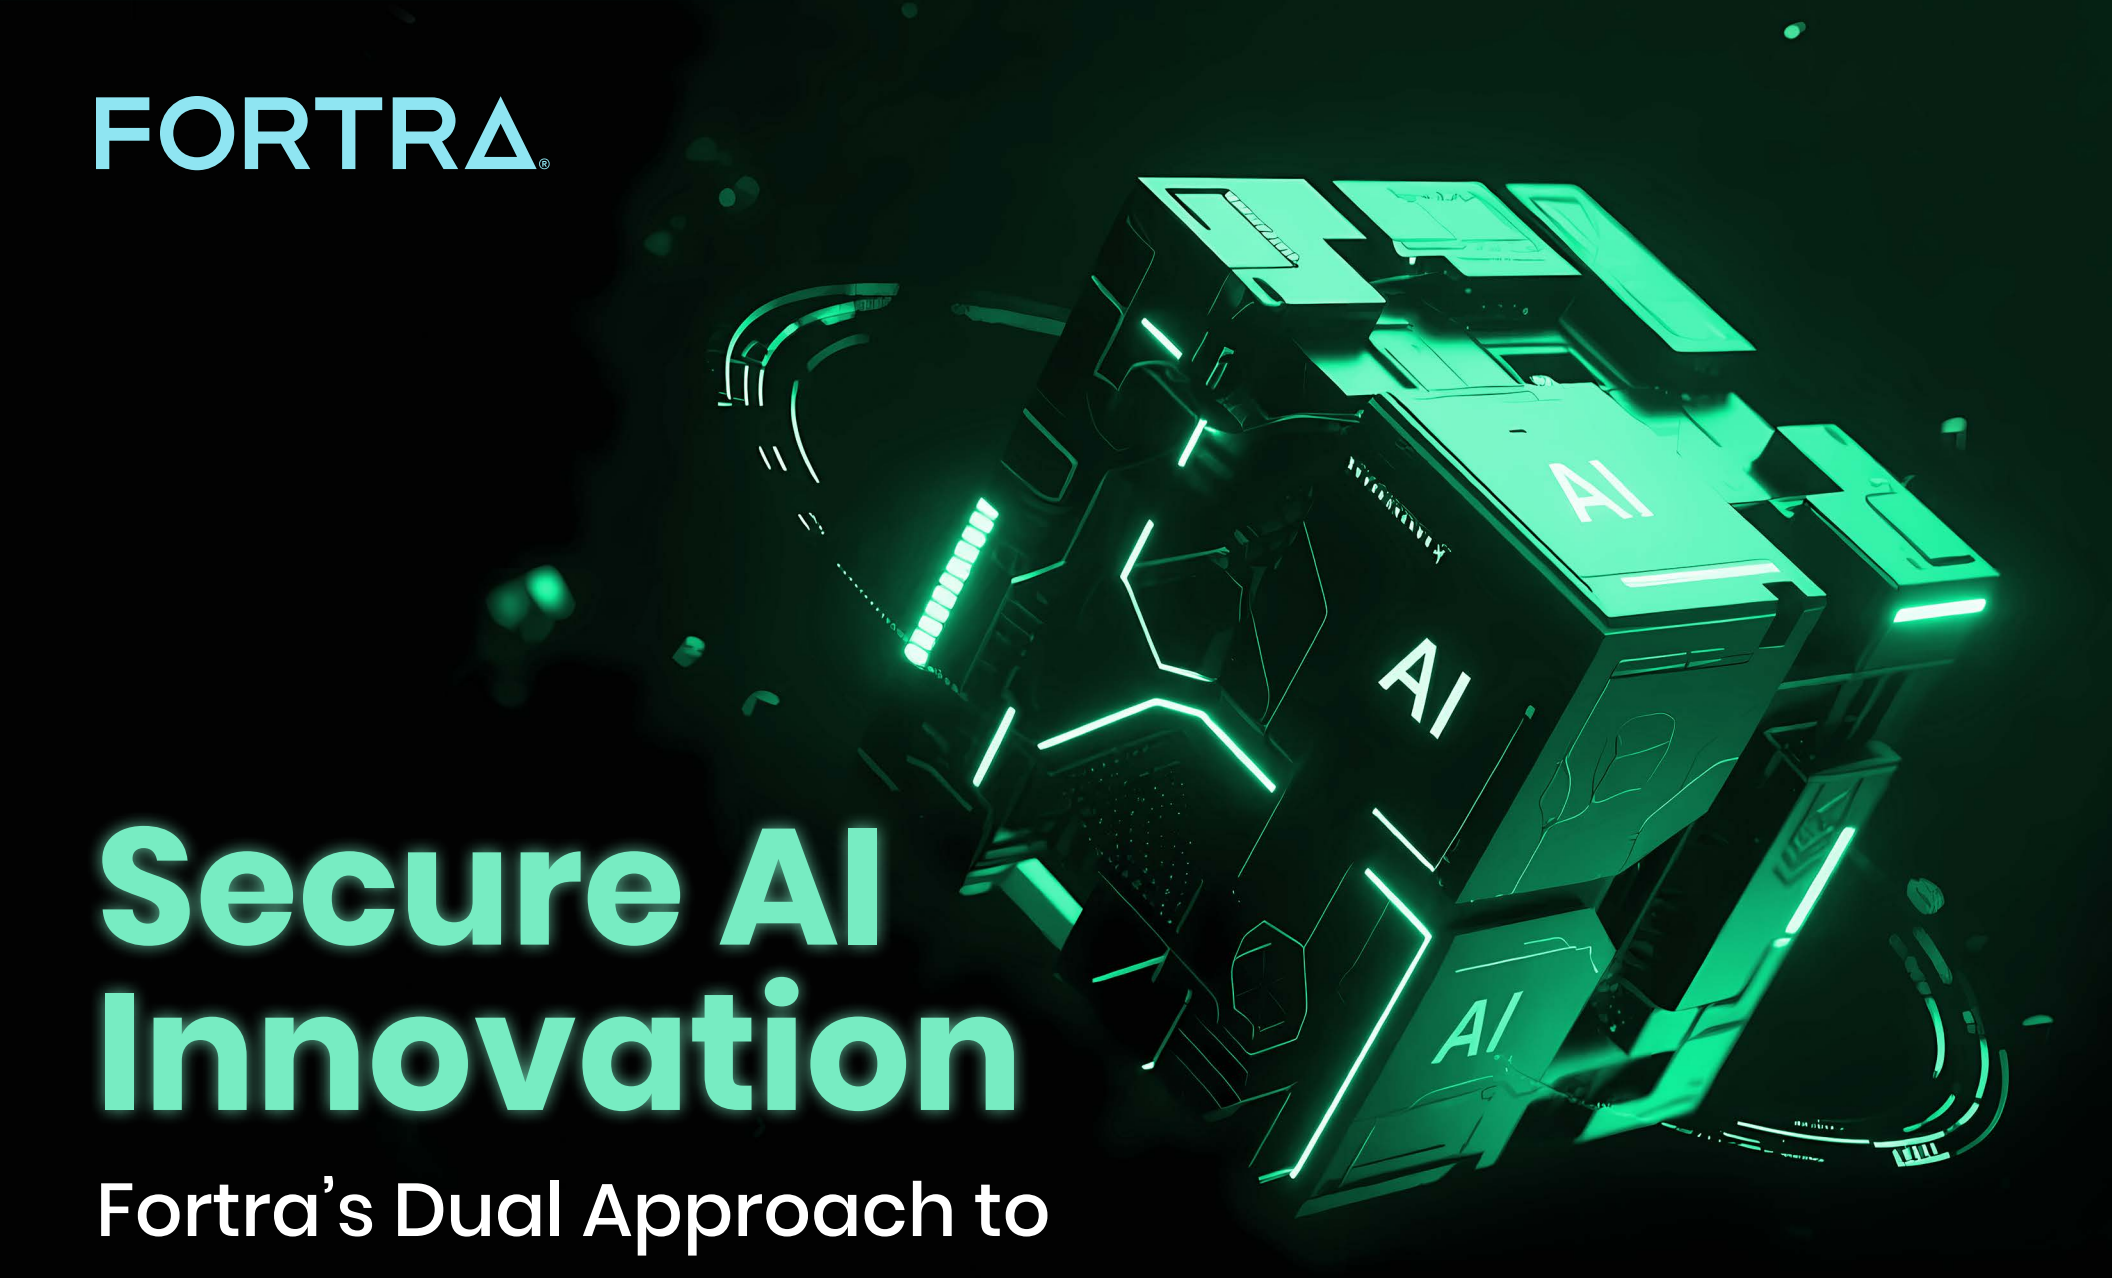

Secure AI Innovation

Fortra's Dual Approach to AI-Driven Cyber Defense and Trusted AI Advancement

Table of Contents

Executive Summary

AI and machine learning are transforming how organizations innovate, enabling faster decision-making, smarter automation, and more personalized customer experiences than ever before. Yet this transformation is only as strong as the data that powers it—and that data faces unprecedented security challenges.

Data is the lifeblood of AI. Without secure, high-quality data, AI systems become vulnerabilities rather than advantages. As organizations rush to implement AI solutions, they often overlook a critical truth: compromised data leads to compromised AI, exposing businesses to new attack vectors including prompt injection, training data poisoning, and AI-powered social engineering campaigns.

Enter Fortra

With over four decades of cybersecurity expertise, Fortra's security experts have witnessed and successfully defended against countless evolving attack surfaces. Each technological advancement—from cloud computing to mobile devices to IoT—brought with its own security challenges that required innovative defense strategies. AI follows this familiar pattern.

The AI arms race is here. Threat actors are rapidly weaponizing AI to launch more sophisticated attacks—generating convincing phishing content, automating vulnerability discovery, and creating deepfakes for social engineering at unprecedented scale and speed. In this accelerated threat landscape, every security company must embrace AI not as an option, but as an operational necessity to detect, analyze, and respond to threats at machine speed.

Our foundation: Secure data is secure AI—and AI-powered defense is essential survival.

Fortra understands that AI security demands a comprehensive approach. We leverage AI's capabilities to outpace threat actors while simultaneously addressing the unique vulnerabilities that AI systems themselves create. Every AI implementation must be built on a bedrock of data security and integrity, with comprehensive protection throughout the entire data lifecycle—from collection and storage to processing and model training.

This whitepaper demonstrates how Fortra's balanced, data-first security approach enables organizations to harness AI's transformative power while maintaining the highest standards of protection against both AI-enhanced cyber threats and AI-specific vulnerabilities.

3 Pillars: Enhancing Security OF, FROM, and WITH AI

AI security requires a comprehensive approach that addresses three distinct but interconnected challenges. Fortra's framework encompasses:

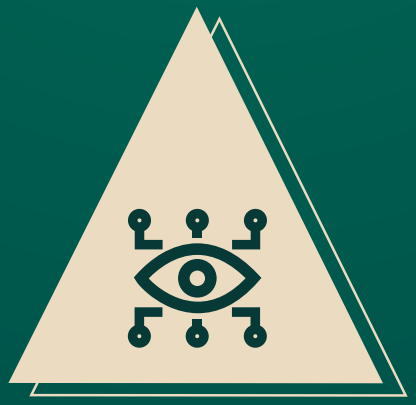

Security OF AI: Protecting AI Systems and Innovation

How Fortra safeguards your AI implementations, models, and data pipelines to ensure secure innovation and maintain system integrity.

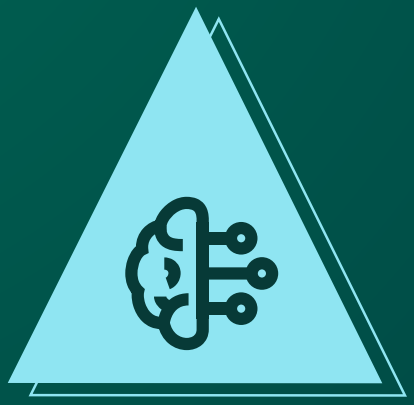

Security FROM AI: Defending Against AI-Enhanced Threats

How Fortra identifies and neutralizes sophisticated attacks that leverage AI capabilities, from AI-generated phishing to automated vulnerability exploitation.

Security WITH AI: Leveraging AI to Transform Business Operations

How Fortra integrates AI technologies to enhance business processes, improve operational efficiency, and drive innovation across the organization while maintaining security standards.

Whether you're building with generative AI or defending against AI-powered attacks, Fortra's holistic three-pillar approach provides a clear path to secure innovation—without compromising trust, compliance, or data quality.

AI Risks and Opportunities

Generative AI is transforming business operations, largely for the better. However, this transformation demands fundamental changes in how we secure data, infrastructure, and access. Organizations face a complex array of challenges as they navigate the emerging GenAI landscape:

Regulatory Compliance and Data Governance

Organizations must demonstrate responsible data usage while establishing and monitoring clear boundaries for AI deployment. This encompasses protecting sensitive content in training datasets, managing exposure during inference operations, and preventing inadvertent data leakage through unauthorized third-party AI services. The challenge intensifies as employees increasingly use shadow AI tools without IT oversight, potentially exposing proprietary data to external models. As regulatory frameworks evolve rapidly, maintaining compliance requires robust data governance, continuous monitoring, and clear policies governing AI tool usage across the organization.

Access Control and Security Gaps

Proper access controls for both users and systems represent the most critical—and most overlooked—aspect of AI security. According to IBM research, 97% of organizations with breached AI models or applications lacked any access controls whatsoever. This stunning statistic reveals a fundamental security blind spot in AI deployments. Like critical infrastructure, GenAI pipelines require comprehensive security frameworks including mandatory identity verification, tiered permissions for different model operations, granular data access controls, and complete audit trails.

Data Integrity and Model Reliability

Ensuring data accuracy for large model training presents ongoing challenges that compound over time. LLM performance, reliability, and trustworthiness depend directly on the quality of input data. Maintaining data integrity at scale protects models from incorrect biases and flawed conclusions, preserving their usefulness and ensuring genuine improvement over time. Poor data quality doesn't just reduce performance—it can embed harmful biases that persist throughout the model's lifecycle, affecting every output and decision the system makes.

These risks are compounded when organizations treat AI systems as static applications rather than dynamic, semi-autonomous agents. Successfully navigating these challenges isn't just about risk mitigation; it's about unlocking AI's transformative potential while maintaining trust, compliance, and security.

Pillar 1: Security OF AI: How Fortra Protects AI and Innovation

Setting the Stage

GenAI systems are only as secure as the infrastructure, data, and access layers that support them. Fortra secures your data wherever it lives—on-premises databases, cloud applications, or AI-native SaaS platforms—ensuring comprehensive protection throughout your entire AI ecosystem. Our integrated stack of capabilities addresses GenAI-specific risks while preserving performance, flexibility, and innovation speed.

Technical Deep Dive

Data Classification and Protection

Fortra's Data Classification and Data Loss Prevention (DLP) solutions label sensitive content, enforce controls, and stop unauthorized transfers from shadow LLMs, where users paste proprietary material into public tools. Whether the threat is intentional exfiltration or accidental upload to a GenAI tool, Fortra helps organizations comply with internal policy and external regulation. Our comprehensive Data Security Posture Management (DSPM) strategy combines classification, DLP, and policy enforcement to help customers discover sensitive data, understand where it resides, and enforce access and sharing rules across hybrid environments and GenAI tools. This enables organizations to reduce shadow AI risk, stop sensitive data leakage, and maintain compliance across the AI lifecycle. Additional features like secure collaboration allow teams to share AI outputs or training datasets safely, with granular permissions and usage logging. Features like secure collaboration allow teams to share AI outputs or training datasets safely, with granular permissions and usage logging.

Providing Visibility into Data Science Infrastructure Vulnerabilities

LLMs and ML systems often rely on a patchwork of orchestration frameworks, open-source libraries, and cloud services—each introducing potential risks. Fortra Vulnerability Management (VM) continuously identifies and prioritizes risk across codebases, data pipelines, and third-party components. Our vulnerability ranking uses supervised learning to enhance CVSS scoring with exploit likelihood, surfacing the issues that matter most for timely response.

Penetration testing from Fortra Core Impact provides an additional layer of assurance. Teams can simulate adversarial behavior across AI-powered services, identifying whether GenAI workloads can withstand abnormal inputs or logic abuse without performance degradation.

Preventing API Attacks on Applications, Including GenAI

APIs are a key interface for GenAI systems, especially when LLMs are integrated into customer-facing applications or accessed via public endpoints. Fortra solutions protect APIs from threats like those listed on the latest OWASP Top 10 for LLMs. Fortra Managed WAF validates and sanitizes data before it is processed, reducing the risk of malformed requests and other input-based attacks that could disrupt GenAI-powered applications.

Detecting and Responding to Cloud Identity and Infrastructure Attacks

As organizations move data science workloads to the cloud, attackers are increasingly targeting the identity layer, abusing overprivileged accounts, misconfigured roles, or insecure tokens to gain lateral access. Fortra Managed XDR provides 24/7 visibility across hybrid and cloud environments, detecting infrastructure misuse before it escalates into model compromise or data theft. This continuous monitoring spans the full attack surface: from IAM policies to suspicious network patterns, abnormal model activity, or unauthorized data egress.

Edge and Access Security for GenAI

At the perimeter, Fortra has integrated multiple technologies to protect AI systems from unauthorized access and input-based exploitation. Our Managed WAF defends against malicious payloads and prompt injection attempts, while Fortra XDR monitors user, network, and application activity for abnormal behavior.

Moreover, Fortra's Secure Access capabilities—including cloud access security broker (CASB) integrations—ensure that only authorized users can interact with GenAI tools, and that third-party apps adhere to policy-based access controls. Together, these capabilities form a layered defense across the edge and access layers of the AI ecosystem.

- **Secure company inboxes against AI-made phishing threats with** Fortra Email Security, especially Fortra's Cloud Email Security Solution.
- **Spot anomalies that indicate malicious behavior in AI applications with** Fortra XDR; also, leverage the Fortra Platform for accelerated end-to-end detection, investigation, and response.

Pillar 2: Security FROM and WITH AI – How Fortra Uses AI to Enhance Cybersecurity

Setting the Foundation

Fortra doesn't just protect AI—we harness it to protect our customers. However, our approach differs fundamentally from that of other security providers who chase AI trends with oversized, costly models that create more problems than they solve.

Strategic AI Integration: Precision Over Power Built on Years of Experience

With over a decade of successfully deploying AI and ML models in production cybersecurity solutions, Fortra understands that effective AI integration requires precision over power. We've been leveraging machine learning across our security portfolio long before the current AI hype cycle, giving us unique insights into where AI delivers genuine value versus where traditional approaches excel.

We integrate AI strategically across every corner of cybersecurity, but only where it creates measurable impact. The secret is deploying AI as a cost-effective tool precisely where it delivers maximum value. For instance, while enterprises generate massive data volumes, using LLM-powered tools to analyze every data point would create prohibitive costs and latency—taking 10 seconds to minutes per analysis and making real-time security unprofitable.

Most alerts and data points require initial filtering to eliminate duplicates and false positives—tasks where our proven traditional advanced tools outperform expensive AI solutions. Fortra's Threat Brain connects IOCs across entire attack chains using shared intelligence and multi-vector telemetry to eliminate false positives. Our Data Classification and DLP solutions enhance alert accuracy, while the Fortra Platform and XDR automatically investigate notifications through automated playbooks, handling routine threat investigation tasks and reducing SOC alert backlogs.

These essential functions operate efficiently without expensive, computing-heavy AI. We reserve AI, ML, and LLMs for scenarios where they're truly necessary—and use them responsibly, informed by years of real-world deployment experience.

Where Fortra separates itself in the AI arms race is through our unique ML models, developed and refined over more than 10 years of production deployment. As veterans in the space, we currently power several solutions with battle-tested ML/AI and are developing advanced models for:

- Email classification and threat detection
- Discovering specific indicators of compromise (IOCs)
- Web page classification and risk assessment
- PowerShell script classification and analysis

Our advancements leverage LLMs to understand contextual components and generate detailed explanations, assisting threat and intelligence teams with deeper insights. Additionally, our Agentic AI initiatives are progressing rapidly, with several projects in active development.

Technical Deep Dive

Modern adversaries are scaling faster and hiding better using AI. Fortra responds in kind, applying ML, AI, and GenAI solutions across our platform to proactively detect, contextualize, and mitigate threats.

AI-Powered Anomaly Detection Across Systems and Applications

Fortra deploys ML to identify behavioral outliers signaling malicious activity, even in zero-day or novel attack scenarios:

- **Geolocation Anomalies**: Detects logins or access attempts from unusual or geographically inconsistent regions.
- **Kerberos Anomalies:** Identifies deviations in authentication flows suggestive of ticket manipulation or lateral movement.
- **Path, Count, and Time-Series Anomalies:** Flags unusual access sequences, excessive query volumes, or out-of-pattern time-based behavior.
- **Clustering for Malware Classification:** Groups behaviorally similar files to detect novel malware families, regardless of signature.

Dynamic Vulnerability Ranking: Fortra enhances traditional CVSS scoring with ML models that assess real-world vulnerability exploitability. Our system ingests CVEs, network configurations, exploit availability, and telemetry to determine exploitation likelihood—prioritizing patches more effectively. Delivered via Fortra Vulnerability Management, this ML-enhanced scoring supports smarter risk mitigation and more agile remediation cycles.

Fortra Threat Brain: AI-Driven Intelligence Engine

Fortra Threat Brain is our AI-driven intelligence engine fueling proactive, real-time cyber defense across the Fortra security ecosystem. Built on machine learning, large language models (LLMs), and dynamic threat correlation, Threat Brain enables solutions including Cloud Email Protection, GoAnywhere MFT, XDR, and Brand Protection to detect, enrich, and respond to today's most advanced threats.

Combating Automation-Based Attacks

Threat Brain continuously ingests and analyzes telemetry from customer environments and external sources. Using behavioral analytics and clustering, it identifies abnormal patterns signaling automation-based attacks:

- **Spam Campaign Identification**
- **Bot and Account Generation Activity**
- **Service Abuse Detection**

Detecting Generative AI-Powered Threats

As sophisticated social engineering campaigns increasingly rely on generative AI to craft persuasive, human-like content, Threat Brain employs advanced classification models and LLMs for real-time detection:

- **AI-Generated Content Analysis**
- **Brand Abuse and Impersonation Detection**
- **Campaign-Level Threat Intelligence**

LLM-Powered Operational Intelligence

Threat Brain integrates multiple LLMs to elevate threat understanding, accelerate response, and support SOC workflows:

- **Context-Aware Email Threat Classification:** Combines LLM reasoning with historical data and intent analysis to classify complex threats like BEC, internal impersonation, and invoice fraud with greater precision.
- **Triage and Threat Research Support:** Delivers summarized threat insights, campaign linkages, and prioritized IOCs through LLM-powered reasoning.
- **Automated SOC Assistant (In Development):** This system extracts risk indicators, queries external sources for enrichment, retrieves historical incidents, and outputs structured reports to accelerate incident response.

Threat Brain delivers more than threat intelligence—it provides actionable insight, real-time enrichment, and AI-powered precision to the Fortra security platform. It serves as the connective tissue behind smarter detection, faster investigation, and stronger defense against the most advanced threats organizations face today.

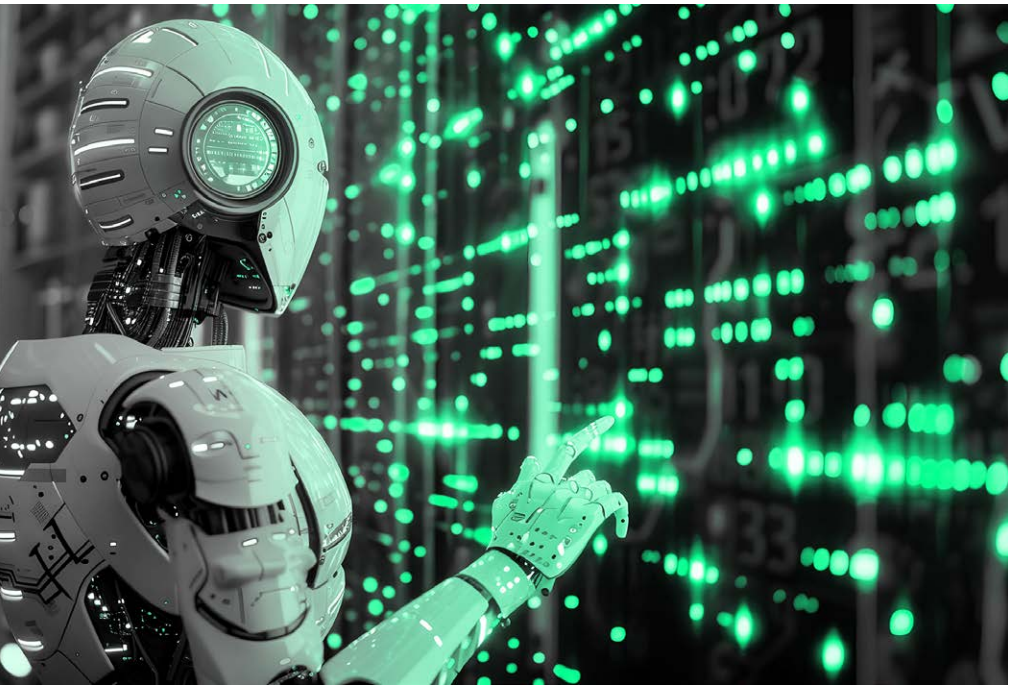

Pillar 3: How Fortra Leverages AI to Transform Business Operations

Leading by Example: Internal AI Innovation

Far from standing on the sidelines of AI innovation, Fortra leverages the same AI technologies and security practices internally that we provide to our clients—ensuring every solution is tested, practical, and proven in real-world use.

We maintain our own team of data scientists developing custom models for customer solutions while simultaneously using third-party AI models secured by Fortra's cybersecurity tools. For example, AI coding assistants used internally are governed by our DLP solutions, addressing the inherent security gaps that exist as AI technology continues evolving.

Internally, Fortra governs AI adoption through a cross-functional Generative AI Council that oversees the rollout of third-party generative tools—from coding assistants and sales copilots to customer support augmentation. Pilots are underway across engineering, marketing, and service teams, each grounded in auditability, human oversight, and measured implementation.

Our focus on AI cybersecurity within internal processes extends beyond preventing data breaches. We believe quality data leads to quality decisions; and you cannot maintain high data quality without robust data security.

Quality Data In, Quality Results Out

GenAI follows the fundamental principle of "garbage in, garbage out." At Fortra, we leverage our advanced cybersecurity capabilities to clean data before feeding AI models—clearing noise, reducing false positives, and removing duplicates. This cleansed, high-quality data produces superior AI outputs and more reliable results.

By applying Fortra's automated alert investigation and advanced filtering to large datasets before model input, we ensure only the most pertinent, accurate information reaches GenAI tools. This approach is particularly critical when using third-party LLMs, where data quality directly impacts output trustworthiness.

The strategy is selective refinement, using only the "best of the best" data to feed models. This approach reduces computing workload while improving accuracy, relevancy, and trustworthiness.

Practicing What We Preach: Firsthand AI Security Innovation

We understand the industry is accelerating toward comprehensive AI adoption, but we advocate for measured, secure progress. The vast data volumes available to organizations can become hindrances without proper vetting, cleaning, and prioritization before feeding GenAI models.

At Fortra, we recognize that data security is fundamental to effective data science. We leverage our cybersecurity expertise to extract optimal information, both for our own operations and for our clients' benefit.

Responsible GenAI Adoption: Accelerating Business Value

Fortra isn't just building AI; we're deploying it responsibly. As in our own environment, we avoid AI hype cycles and focus on customer needs by applying AI precisely where it delivers maximum value.

Our approach is strategic and cost-effective: comprehensive security solutions handle routine tasks and data preparation, while powerful LLM models address complex challenges. Only the "tough cases" warrant LLM investigation—validated anomalies, confirmed threats, and IOCs that could reveal advanced persistent threats or emerging malware.

By creating targeted, efficient LLMs that activate when needed, we deliver superior AI benefits while saving clients resources. Our approach provides more value with greater savings, reduced latency, and less operational disruption than security-wide AI overhauls.

Our responsible AI adoption ensures transparent, trustworthy solutions that produce accurate, reliable outcomes our clients can depend on.

Accuracy and trustworthiness are our business value differentiators.

Agentic AI: The Next Governance Challenge

What is Agentic AI?

Agentic AI refers to autonomous systems that can initiate tasks without direct human intervention, leveraging automated reasoning capabilities to make guided decisions on how to solve problems. These systems function similarly to junior employees embedded within an organization's digital environment, performing tasks based on assigned objectives, contextual cues, and learned behaviors.

Why Does It Matter?

Unlike traditional applications, agentic AI can behave in adaptive and unpredictable ways. Its autonomy introduces unique risks that standard governance models cannot manage. Treating these systems as "just another app" underestimates its capacity to evolve behavior based on feedback loops, data exposure, or experimental cues. This leads to outcomes that can diverge from intended objectives.

The Intern Analogy

Managing agentic AI is best approached through a model akin to managing interns. Like interns, these systems require defined identities, role-based access, behavioral expectations, and ongoing supervision. Without clear boundaries and active oversight, they can generate unintended consequences that are difficult to trace or control.

Organizations should already enforce policies around access control, task scope, and performance oversight. These existing mechanisms should be extended and digitized to govern AI agents. This includes implementing programmatic guardrails, robust logging, human-in-the-loop reviews, and regular checkpoints.

A tiered control structure is also critical: narrowly scoped agents handle specific tasks, intermediate agents synthesize outputs from these task-based agents, and supervisory agents oversee overall activity and alignment. This layered architecture supports transparency, reduces risk, and ensures scalable oversight as agentic systems become more integrated across workflows.

Governance Framework: Monitoring, Identity, and Guardian Agents

To govern autonomous AI systems effectively, Fortra applies a three-part model: monitoring agent behavior, enforcing scoped identities, and supervising through policy-aware guardian agents. This framework translates directly into the following best practices:

- **Identity and Role Management:** Assign unique, scoped identities to all AI agents. Ensure roles are enforced with clear boundaries and logging.
- **Continuous Behavior Monitoring:** Use real-time monitoring and anomaly detection to observe AI activity, flag deviations, and initiate reviews:
- **Enforce Access Controls:** Apply least-privilege access across all systems and data touchpoints involved in agentic operations.
- **Adopt Human-Centric Security Practices:** Maintain human oversight with explainability, visibility, and intervention tools embedded into workflows.
- **Design for Privacy and Compliance First:** Build privacy, data protection, and regulatory alignment into the architecture.

Practical Framework for Agentic AI Governance

To bring this model to life, organizations should operationalize governance through repeatable, testable processes:

- **Formal Onboarding/Offboarding of Agentic Systems:** Register agents with unique IDs, roles, scopes, and expiration conditions - just like human accounts.
- **Implement Security and Monitoring Protocols:** Apply continuous monitoring, alert thresholds, and policy enforcement from deployment onward.
- **Run Regular Audits and Behavioral Review Cycles:** Periodically review agent activity against expected behavior, role boundaries, and business outcomes.
- **Align with Human-Centric Security Principles:** Ensure operators can interpret, intervene, and redirect AI behavior with confidence and authority.
- **Incorporate DevOps and Test-Driven Practices:** Deploy agents using the same rigor applied to critical systems: routine testing, version control, and validation that outputs to standardized prompts remain consistent over time.

Together, these frameworks provide a structured, actionable path for organizations to deploy Agentic AI without compromising safety, compliance, or trust. Fortra's solutions – spanning detection, classification, monitoring, and access control – serve as the foundation for turning policy into practice.

Agentic AI: Real-World Use Cases & Lessons Learned

Agentic AI is rapidly moving from pilot to production. A Capgemini study projects that AI agents could generate as much as $450 billion in value by 2028, yet only 2% organizations have scaled deployments. Another 47% are in active pilots – mostly in customer service, finance, and healthcare.

In these sectors, agents are automating tasks like ticket triage, fraud detection, and clinical documentation, but autonomy alone isn't driving results. Success relies on governance: role enforcement, real-time monitoring, scoped permissions, and human oversight.

A Wolters Kluwer study found that 44% of finance leaders expect to deploy agentic AI within a year. In a separate PagerDuty study, 62% of organizations using agents expect to achieve a triple-digit ROI; this underscores both the opportunity and the need for governance.

Fortra's governance model supports this shift, enabling customers to scale agentic AI securely and accountably.

Conclusion: Innovation Without Compromise

In a climate where organizations are quick to fall for the hype of ubiquitous AI adoption Fortra likes to approach the issue with a level head.

Is AI essential to securing modern architectures, both now and in the future? Considering the threat of AI-powered exploits, absolutely. But is painting everything with a broad AI brush the answer? Unless companies have unlimited time and financial resources, no.

At Fortra, we know our customers personally and understand the challenges they face. Our purpose is to develop AI-infused solutions that pinpoint problems and solve them with the technology that's best suited to the job. Sometimes that's Agentic AI, sometimes Generative AI, and sometimes it's a powerful Large Language Model.

Either way, Fortra secures AI innovation so you can confidently leverage the tools that powerful technology providers, like AWS, have to offer.

Through our own internal experience and the industry expertise that comes from over four decades of continuous cybersecurity protection, we can provide our clients with the customized protection they need in the era of artificial intelligence.

That means security for internal AI models and against malicious AI models, allowing companies to grow with the AI tools and approaches that meet their independent needs. Most importantly, we do it with the know-how of a company that uses, invents, and innovates AI models of our own.

Protect Your Data in the Age of AI

The foundation of AI is data. Start with your data. Know where it lives, how it moves, and who can access it.

See how Fortra's DLP and classification solutions identify, monitor, and secure sensitive information across GenAI use cases. Then, lean on Fortra to build with confidence and innovate responsibly.

LEARN MORE

About Fortra

Fortra provides advanced offensive and defensive security solutions that deliver comprehensive protection across the cyber kill chain. With complete visibility across the attack chain, access to threat intelligence spanning the globe, and flexible solution delivery, Fortra customers can anticipate criminal behavior and strengthen their defenses in real time. Break the chain at fortra.com.

Copyright © Fortra, LLC and its group of companies. Fortra®, the Fortra® logos, and other identified marks are proprietary trademarks of Fortra, LLC.

fta-gd-0925-r2-ak