Confidently harnessing AI:

Establishing your AI governance and security framework

Table of contents:

Introduction: Navigating the AI tightrope – innovation vs. risk

The rapid rise of AI has led to urgent new realities and a core dilemma for business leaders: fear of missing out on trillions of dollars of potential value versus fear of getting it wrong. Research suggests that generative AI will add trillions annually to the global economy in the next few years, even as headlines trumpet data breaches, hallucinated content, copyright lawsuits, and regulatory fines. Meanwhile, governance remains a challenge. Notably, PwC found that "50% of executives cite translating Responsible AI principles into operational processes as their biggest barrier to progress." [1]

But let's be frank about that adoption barrier: AI has already infiltrated your organization, creating unmanaged "shadow AI" risks beneath the surface. Marketing is drafting content with consumer chatbots, engineers are generating code with AI assistants, and your vendors are sending your proprietary data to cloud-hosted third-party models.

This makes the AI governance challenge — to prevent risk and enable safe innovation — a leadership imperative. "It's essential that business leaders install guardrails, implement security frameworks, and establish effective governance structures. Effective governance isn't just about stopping risk; done right, it enables safe innovation, so your teams can confidently leverage AI as a powerful tool," notes Aaron Richmond, senior solutions architect and AI governance expert at SHI.

[1]PwC, "2025 Responsible AI survey", October 2025

Essential guardrails: Defining the boundaries for safe innovation

You've processed waves of change before — from data integration with suppliers to social networks and financial crises. As business and IT leaders, your role is to ensure innovation can move fast without creating unnecessary risk. That means understanding emerging threats and applying policy, education, monitoring, and technology to mitigate them.

Protecting enterprise data: Your most critical asset in the AI era

The most essential AI guardrail is data protection. When teams use AI tools to analyze customer data, process financial information, or generate content based on proprietary knowledge, they may, often unintentionally, send cloud AI models vast amounts of sensitive information. This could prove catastrophic if mishandled. Beyond accidental leaks, unsecured data streams are targets for malicious actors, and feeding proprietary data to public LLMs can inadvertently train models that benefit your competitors, subtly eroding your market advantage.

You will have to decide when to use cloud models (such as those from OpenAI and Anthropic) and when to use private models. The tradeoffs are substantial. Cloud model vendors may use your company's data to train their models — essentially storing your data in a form that could be queried through the vendors' AI models, including private data and intellectual property you did not intend to share.

While private models that run within the business' security boundary offer control, they may lack the cutting-edge capabilities or cost-effectiveness of cloud options, unless the business makes a significant investment and brings in specialized expertise. Today, the most advanced models are typically delivered by cloud providers. Open-source models are closing this gap, but they require considerable scale and sophistication to be truly enterprise-ready and economical.

Given this, organizations should begin with a pragmatic, risk-based assessment of AI vendor guarantees against their own requirements. Shared responsibility frameworks can clarify provider commitments, but these must be weighed against regulatory and risk obligations to identify residual gaps the organization must address, such as deploying data loss prevention (DLP) tools to protect sensitive information.

“Data is the lifeblood of AI, but, uncontrolled, it’s a critical liability,” notes Cory Peters, vice president of product and MSP, SHI. “Visibility into this data flow is the first step toward establishing essential control.”

SHI helps organizations establish critical guardrails through policy development, employee education, monitoring, and technology implementation. In parallel, SHI assesses AI vendor guarantees against each organization's risk tolerance, regulatory requirements, and specific business needs.

Assess your AI readiness

Managing privacy and compliance: Navigating a tangled global web

AI systems introduce specific privacy considerations that directly impact legal compliance with a growing, tangled web of regulations like GDPR, CCPA, and newer AI-specific regulations such as the European Union’s AI Act. The cross-border nature of many AI implementations creates complex regulatory challenges, potentially exposing businesses to significant liability and reputational damage.

Organizations must establish clear policies and actionable guidelines — not just shelfware — that define permissible data sharing with AI systems, access controls for AI-generated outputs, customer consent protocols, data retention limits, and deletion requirements. The urgency is growing. According to Gartner®, "By 2027, more than 40% of AI-related data breaches will be caused by the improper use of generative AI (GenAI) across borders.”[2] Yet operationalizing these requirements remains difficult. “Translating complex global privacy laws into practical AI controls is daunting,” explains John Moran, SHI’s director of AI product management. “Companies struggle to demystify the requirements and build privacy directly into their AI workflows, which is a must for fostering trust and ensuring compliance.”

[2] Gartner®, “Gartner Predicts 40% of AI Data Breaches Will Arise from Cross-Border GenAI Misuse by 2027,” 17 February 2025

Controlling AI tool usage:
Taming the double-edged sword

The proliferation of AI tools is a double-edged sword, boosting productivity while introducing unmanaged risks through "shadow AI." Clear usage guidelines for employees are essential. Companies must define which AI tools and platforms are approved for business use and how they can be applied, balancing security with business agility.

Maintaining an effective "allowlist" requires a defined process as employees clamor to get the benefits of new tools. As SHI's Richmond observes, "Most employees are not malicious — they just need a guide to doing the right thing, with guardrails to prevent missteps. Successful organizations will translate their policy and solution strategy into guides and guardrails."

SHI's AI Security Posture Review and AI Security Briefing are complimentary offerings designed to guide organizations through the secure implementation and ongoing protection of AI applications and the security tools needed.

Enterprise versions of tools such as ChatGPT, Microsoft Copilot, Glean, and Google Gemini offer greater control over data handling than their consumer counterparts. These enterprise solutions provide access controls and data protection features needed to manage risk effectively while enabling productivity.

Managing AI-generated content:
Ensuring quality, accuracy, and accountability

AI-generated content can deliver real productivity gains, but it also creates new risks related to ownership, attribution, accuracy, and sound business judgment, which blur the lines of responsibility. Organizations need clear policies on ownership of AI-generated work, disclosure of AI involvement, and verification before use. Who is liable if AI hallucinates plausible but incorrect advice?

Collaborating effectively with AI demands new skills and ways of approaching work. Instilling these in your organization requires both active education and clear policy.

Leveraging the SHI AI Readiness Workshop will help with evaluating your organization's level of AI skills and education.

"AI-generated content requires a new level of scrutiny," observes SHI's Moran. "Companies need to set up review workflows, implement AI detection where appropriate, and train teams on responsible AI co-creation to maintain quality and trust."

Addressing ethical considerations:
Building AI on a bedrock of trust

AI systems can inadvertently perpetuate biases that expose businesses to reputation damage, legal liability, and erosion of customer and employee trust. When AI influences decisions about customers, employees, or partners, ethical missteps can quickly become business disasters.

As SHI's Richmond sees it, "Ethical and responsible AI can be summed up in four questions. Ask yourself: "What is right and wrong in our organization? What rules apply to us? Do we know what we are responsible and accountable for? Can we prove that we are meeting these responsibilities?"

The answers will help organizations develop clear, actionable, and ethical principles for AI use that address fairness, transparency, and accountability. Microsoft's Responsible AI Standard[3] offers a useful framework for this purpose that emphasizes fair treatment across groups.

Implementing algorithmic fairness audits for high-stakes domains helps prevent discriminatory outcomes that could violate laws and damage relationships. Many organizations establish ethics review boards with diverse representation to assess systems against guidelines and to provide guidance for complex cases.

SHI provides frameworks and best practices for structuring these effectively, ensuring diverse representation and clear mandates to navigate complex ethical dilemmas.

Strengthen your AI governance with SHI

[3] Microsoft, "Microsoft's framework for building AI systems responsibly," June 2022

Security framework: Where AI policy meets technical reality

A comprehensive security framework provides the mechanisms and processes to protect AI systems and the data they handle. Such frameworks translate guardrails into a technical implementation that safeguards business assets, customer data, and brand reputation.

Threat assessment and risk prioritization: Illuminating AI's blind spots

Traditional security assessments often miss AI-specific risks such as model inversion attacks, prompt injection, and data poisoning. Are your defenses prepared for these?

Organizations must map data flows between AI components and existing systems, identify potential attack vectors, and prioritize risks based on business impact. "Red team" exercises, where security experts simulate attacks on AI models, help uncover vulnerabilities before they can be exploited in production.

Allowing AI to write, deploy, and execute code: Balancing velocity with vigilance

A software development team that integrates AI into daily work can deliver some of the largest gains from today's AI — doubling software engineering productivity is realistic. But it can also create significant risks if not managed carefully. Crisp boundaries are essential. Some can be enforced by software, but most can only be put into practice through clear policy, good education, and oversight. Specific risks include deploying AI-written code without sufficient review and allowing AI agents to execute code on their own initiative.

Implementing technical safeguards: Building a layered defense

Organizations must block employee and system access to unauthorized AI tools and platforms using approaches such as cloud access security brokers (CASB) or secure web gateways (SWG). Then, enforce data usage policies with technical safeguards to the extent feasible.

Beyond standard data protection, AI systems benefit from a layered defense including specialized safeguards, such as federated learning and differential privacy, which maintain data utility while enhancing protection. Organizations operating across jurisdictions should implement data localization strategies to comply with sovereignty laws. Instead of relying on freeform chat, design constrained interactions to minimize the information that is passed to the AI application.

SHI helps organizations protect AI systems and data through a comprehensive security framework, threat modeling, and technical safeguards, while enabling secure AI-assisted coding practices and specialized monitoring for effective incident response.

Enhance your cybersecurity with SHI's AI Security Posture Review

[4] EY, "Cyber and AI oversight disclosures: what companies shared in 2025", October 2025

Monitoring and incident response: Maintaining a vigilant watch

AI systems require specialized monitoring that tracks both technical metrics (like unusual query patterns) and business outcomes (like customer complaints). Based on these signals, organizations should establish AI-specific incident response plans that address model manipulation, data poisoning, and inappropriate outputs. Organizations should also create dedicated response teams, combining data science and cybersecurity expertise, and conduct regular tabletop exercises to prepare for novel threats that traditional security teams might miss.

Security testing and validation: Battle testing your AI defenses

Regular security testing is crucial to battle-test your defenses before attackers can exploit them. "AI amplifies existing weaknesses in security posture, particularly concerning data and systems access," explains Christina Halikiopoulos, AI security solutions architect at SHI. "It's essential for companies to establish a robust security foundation before deploying AI solutions, and to implement technical controls to safeguard these systems during runtime." Companies should implement comprehensive testing protocols, including penetration testing, adversarial testing, and sensitivity analysis. EY research found that 58% of companies they surveyed in 2025 report that "their cybersecurity preparedness includes simulations, tabletop exercises or response readiness tests."[4]

Policy and governance: Bringing your AI strategy to life

Effective governance transforms AI guardrails and security frameworks from concepts into operational reality. It's the connective tissue bringing your AI strategy to life, ensuring it's both ambitious and responsible. Without practical governance structures, even the best security measures remain disconnected from day-to-day operations.

"Effective governance isn't about hitting the brakes on AI," Peters says. "It's about building the confidence and clarity to accelerate safely in the right direction, maximizing value while minimizing risk."

Balancing innovation and risk management:
The delicate dance

AI governance must balance enabling innovation with managing risk. Organizations should establish clear principles for responsible AI use that reflect their specific values and risk tolerance, and then weave these principles into their corporate DNA. The most effective governance committees include legal, finance, IT, security, business units, and executive leadership representatives, ensuring policies address multiple perspectives and remain practical for real-world application.

Aligning AI with business:
Ensuring value, not just compliance

Governance must directly address core business objectives alongside risk mitigation. Organizations should establish processes for evaluating AI initiatives against strategic goals, while also calculating a targeted return on investment. EY's 2025 Responsible AI Pulse survey found that among organizations that have adopted Responsible AI, "eight in 10 respondents report improvements in efficiency and productivity" while "about three in four say it has improved their ability to understand customers and respond quickly to shifting market conditions."[5]

[5] EY, "Responsible AI Pulse survey," October 2025

Establishing governance structures:
Creating workable frameworks

Establish an AI governance committee, with clear authority and resources, to develop policies, evaluate high-risk applications, and ensure alignment with values. This committee is responsible for creating clear decision ownership and escalation paths based on risk levels, preventing governance from becoming a bottleneck while ensuring appropriate oversight where needed. Forrester's version of the RACI matrix for AI provides a useful template for governance, designating specific roles as responsible for model fairness, legal teams as accountable for compliance, and executive leadership as ultimately accountable for overall governance.[6]

Implementing continuous improvement:
A living, breathing framework

AI governance must be a living, breathing framework that evolves with technology. Organizations should implement automated compliance checks where possible and establish regular review cycles that evaluate both technical performance and alignment with governance principles. "If you're truly committed to governance and risk management, you'll be willing to do the work. It's not trivial, but it is possible," notes Richmond. This commitment to ongoing monitoring distinguishes mature governance approaches from superficial compliance frameworks.

"We help organizations test the integrity of their AI systems… either in real-time or after an event, to catch anomalies traditional tools might miss."
- Lee Ziliak, Field Chief Technology Officer at SHI

Empower your AI strategy with SHI

[6] Forrester, "The AI Governance RACI Matrix," 2024

Outcomes:
Enabling safe, scalable AI adoption

Organizations that establish strong AI governance and security frameworks move from unmanaged experimentation to confident, responsible adoption. Clear guardrails spanning policy, education, monitoring, and technical controls, enable teams to innovate safely while reducing shadow AI risk. Employees gain clarity on approved tools, acceptable data usage, and when human oversight is required, building trust and accelerating adoption across the business. This foundation sets the stage for meaningful, repeatable business impact.

Business impact: Protecting value while unlocking innovation

Effective AI governance protects sensitive data, intellectual property, and brand reputation while enabling AI to be applied in high-value workflows. Organizations reduce the risk of breaches, compliance failures, and reputational damage, while ensuring AI initiatives align to business outcomes and ROI. When governance is operationalized, rather than treated as shelfware, it becomes an enabler of speed, not a constraint, allowing innovation to scale with confidence. Defined accountability, continuous monitoring, and incident response further mature the organization, turning risk management into a durable, evolving capability.

Why it matters now:

AI is already embedded in daily work, often outside formal controls. Without governance, risk grows faster than value. With it, organizations replace uncertainty with confidence and enable AI to scale responsibly.

Bottom line:

Strong AI governance turns AI from a source of risk into a trusted, scalable business capability, enabling innovation today while protecting the enterprise for what comes next.

Next steps — and how SHI can help

Embarking on AI governance can feel overwhelming, but SHI provides a clear path forward. Focus on practical actions delivering immediate value. SHI's AI Readiness Workshop offers a structured approach, meeting you where you are and guiding you toward comprehensive, confident AI adoption. Here's how SHI helps you take those crucial first steps:

Step 1

Develop an AI governance charter. SHI facilitates focused executive workshops to secure senior sponsorship that clarifies roles, responsibilities, and decision rights — laying the foundation for your AI governance

Step 2

Implement a use case capture system. SHI works with you to establish and enable your AI Council, implement a use case capture system, and develop a mature ROI analysis process to support innovation.

Step 3

Establish a data strategy. SHI's experts work with your teams to create a data governance plan that enables innovation while protecting sensitive information.

Step 4

Create an AI literacy program. **SHI can help you** implement frameworks like the META AI Literacy Scale to ensure employees understand safe and effective AI use.[7]

Step 5

Inventory current AI usage and implement technical guardrails. We use specialized tools and our own proprietary methodology to help you uncover both sanctioned and shadow AI, across all departments, to identify any existing governance gaps. SHI helps select and implement essential controls around approved AI tools, including DLP and access management, that balance security with productivity.

[7] SHI AI Literacy Blog, "Here's why AI literacy is easily the best thing you're not doing," 2024

Select, deploy, and manage technology

SHI also offers expertise and implementation capability to support an enterprise's entire AI journey. This can begin with an AI Readiness Workshop that evaluates an organization's security posture, data readiness, and infrastructure capabilities. The workshop brings together key stakeholders to develop tailored frameworks that align with specific business objectives. "Most organizations stumble on where to start," says Richmond. "They need help writing policies, putting together comprehensive plans, and selecting the right vendors for their needs."

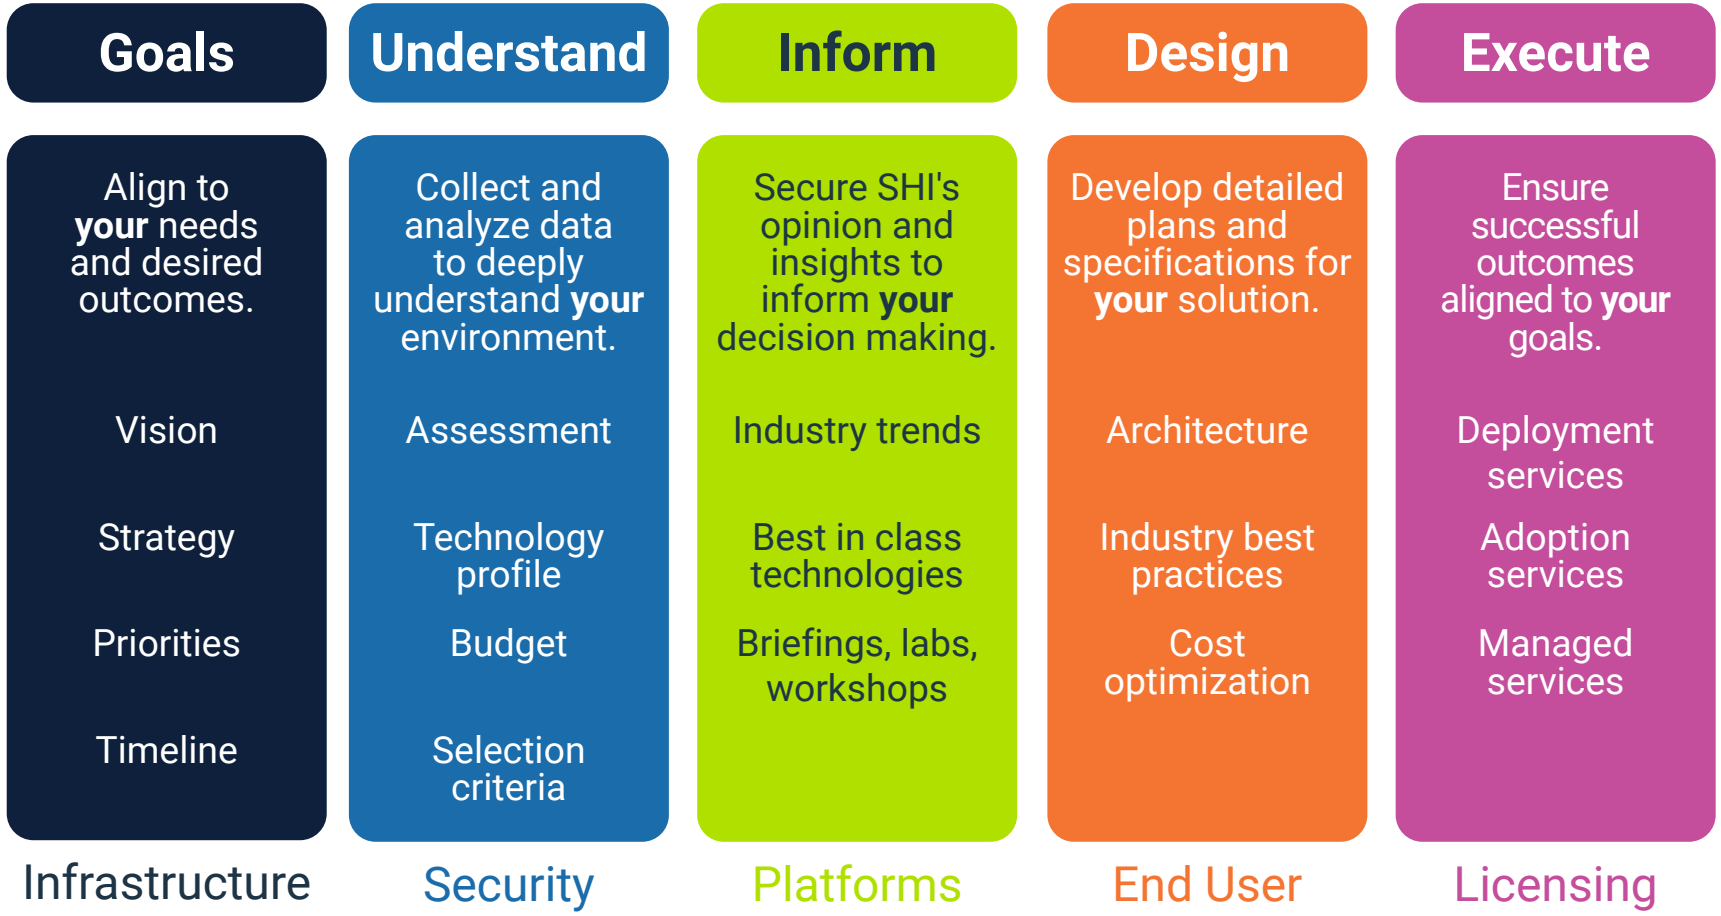

These fundamentals provide a strong starting point for organizations. Turning these into a fully integrated, enterprise-grade AI capability requires deeper assessment and hands-on execution.

SHI works with organizations to move beyond concepts and point solutions, delivering end-to-end advisory and implementation that connects people, process, technology, and governance. Through structured workshops, technical assessments, and ongoing guidance, SHI helps organizations operationalize AI at scale - aligning strategy to outcomes, managing risk, and accelerating value in a way that is secure, repeatable, and built to evolve.

Additional guidance can be drawn from established frameworks, including:

- NIST AI Risk Management Framework (AI RMF)
- MITRE Adversarial Threat Landscape for Artificial Intelligence Systems (ATLAS)
- Open Worldwide Application Security Project (OWASP) Top 10 LLM Applications 2025
- Google Secure AI Framework (SAIF)

To dig deeper, SHI's AI & Cyber Labs are an excellent resource.

About SHI's AI Security Briefing

SHI is not just a global IT solutions provider; we are your AI transformation partner. We help you navigate complexity, reduce risk, and unlock real business value.

SHI's AI Security Briefing delivers a clear, comprehensive overview of AI security — its importance, common risks, and the strategies and technologies required to safeguard AI models, applications, and data. Participants gain practical insights into best practices for securing AI deployments, identifying vulnerabilities, and implementing robust security measures.

SHI's AI Security Posture and Readiness Review is a complimentary service designed to help organizations securely implement and protect AI applications. This offering provides expert guidance on security strategy at any stage of your AI journey. SHI conducts a detailed evaluation of your current security measures and delivers actionable recommendations to strengthen the protection of your AI applications and the supporting security tools.

Take your AI journey with us.

Connect with an SHI AI expert.