Guidance for Boards of Directors on How to Address AI Risk

A Perspectives on Security for the Board Special Publication

Unleashing the potential of artificial intelligence (AI) promises undeniable benefits. Yet, AI's rapid advancement also carries significant risks. Data breaches, privacy violations, and the potential for biased or harmful outputs present direct threats to an organization's reputation, financial stability, and legal compliance. While there are legitimate concerns around AI, the potential risks of not adopting and utilizing it for business improvement could be far more significant, including the risk of employees using unauthorized AI tools.

Boards of Directors hold a fiduciary responsibility to safeguard their organizations from these intersecting risks. To do so effectively, proactive leadership and strategic oversight of AI initiatives are essential. To help organizations navigate these challenges, we've outlined best practices to streamline and operationalize AI implementation at scale. Additional resources can also be found on our Board of Directors Insights Hub.

Boards of Directors should leverage the following questions and considerations when speaking with your CIO/CTO and CISO on AI related risks.

Stakeholder Identification

- **Ask**: Is there a multidisciplinary team of experts assembled to assess AI initiatives before deployment and provide ongoing evaluation?
- **Consider**: Typically this would include representatives from various functions: IT Infrastructure, Information Security, Application Security, Risk, Compliance, Privacy, Legal, Data Science, Data Governance, and Third-Party Risk Management teams.

Oversight and Escalation Channels

- **Ask**: Have we established points of escalation so that when questions arise, there's a clear path to getting them answered?
- **Consider**: While some organizations may want to send their internal queries to Legal, Compliance, or Information Security, others might prefer to empower a designated committee to make decisions. Implement mechanisms for providing visibility on the status of each AI initiative both to internal stakeholders.

(Guidance for Boards of Directors on How to Address AI Risk, cont'd.)

Guiding Principles

- **Ask**: Have we defined our organization's guiding AI principles to articulate foundational requirements and expectations, as well as use cases that are explicitly out of scope?
- **Consider**: Guiding principles should be flexible and not overly prescriptive, capturing commitments from which the organization won't deviate; for instance, a focus on safeguarding customer privacy, or ensuring a human is involved in reviewing AI-generated decision making for certain use cases. As an example, see Google's Responsible AI Principles.

AI Framework

- **Ask**: Are we using a framework [such as Google's Secure AI Framework (SAIF)] for a secure and consistent approach to AI implementations?
- **Consider**: Beware of security framework traps — a framework is just a tool, and its use shouldn't be confused with having achieved your objective. Rather, a framework like SAIF is a helpful way to approach AI implementation to ensure its multiple facets are comprehensively considered.

Policies and Standards

- **Ask**: Have we documented and implemented relevant policies and procedures for AI design, development, deployment and operations?
- **Consider**: Ownership of these resources typically varies by team, and effective AI governance oversight requires a concerted effort to maintain accuracy, completeness and alignment.

Risk Appetite

- **Ask**: How well do our current security controls address the specific risks of AI adoption, and does this align with the Board's established risk appetite?
- **Consider**: Rank use cases in order of business priority and the degree of risk they may pose, tailoring the security and data protection controls accordingly.

Data Governance

- **Ask**: Are we considering our organization's data governance program in our AI strategy, as AI models typically require high-quality data that should be appropriately sourced, collected, cleansed, stored, protected and normalized in compliance with data privacy regulations?
- **Consider**: Integration with the data governance program is crucial for building trustworthy AI systems, ensuring regulatory compliance, and maximizing the value derived from AI initiatives. Carefully selecting your data set and tuning the AI model for your specific needs can also help minimize the potential for hallucinations and risk of prompt injections.

Boards of Directors play a pivotal role in shaping the future of AI within their organizations. By embracing proactive oversight and adopting best practices, boards can harness the transformative power of AI while mitigating its inherent risks. This responsible approach will ultimately build trust and ensure the long-term success of AI-driven initiatives.