Workshop #4:

From the Experts

Balancing AI Innovation with Data Security and Governance

DATA SECURITY CREATIVE COUNCIL

DATA SECURITY
CREATIVE COUNCIL

Contributing Members

Sonali Bhagwat
SR. Director of Data Governance
Adobe

Venkat Valleru
Principal Information Security and Compliance Engineer
Informatica

Jeff Farinich
CISO
New American Funding

Archit Uppot
Senior Product Manager
Netskope

Arthur Hedge
President
Castle Ventures

Karen Lopez
Data Evangelist
InfoAdvisors

Workshop #4:

From the Experts

Balancing AI Innovation with Data Security and Governance

Organizations are grappling with effectively balancing AI's transformative potential with the equally crucial data security and governance needs. This white paper is a direct and timely response to this pressing issue, stemming from a workshop where leaders in data and security, such as **Sonali Bhagwat** (SR. Director of Data Governance at Adobe), **Jeff Farinich** (CISO of New American Funding), **Arthur Hedge** (President of Castle Ventures), **Karen Lopez** (Data Evangelist at InfoAdvisors, and **Archit Uppot** (Senior Product Manager at Netskope) convened to discuss these critical matters. Their insights form the basis of exploring the complexities of integrating AI into enterprise operations.

The discussion highlighted three primary concentric circles of AI adoption:

- **Companies using AI tools and services**
- **Those building AI applications for specific use cases**
- **AI-first companies where AI is central to the business model.**

Each circle presents unique opportunities and risks, particularly regarding data security.

Key takeaways from the workshop include the importance of data governance, transparency, and the clear labeling of AI-influenced decisions. These are crucial in addressing concerns such as the potential for cross-pollination of sensitive data, unauthorized access through embedded AI tools, and ensuring compliance with data protection laws when integrating AI into the web of business infrastructure.

This white paper also explores the development and implementation of AI policies, addressing the challenges posed by AI's novelty and the lack of internal legal expertise. It provides practical recommendations for organizations looking to leverage AI while safeguarding critical data, ensuring that AI's benefits are realized without compromising security and integrity.

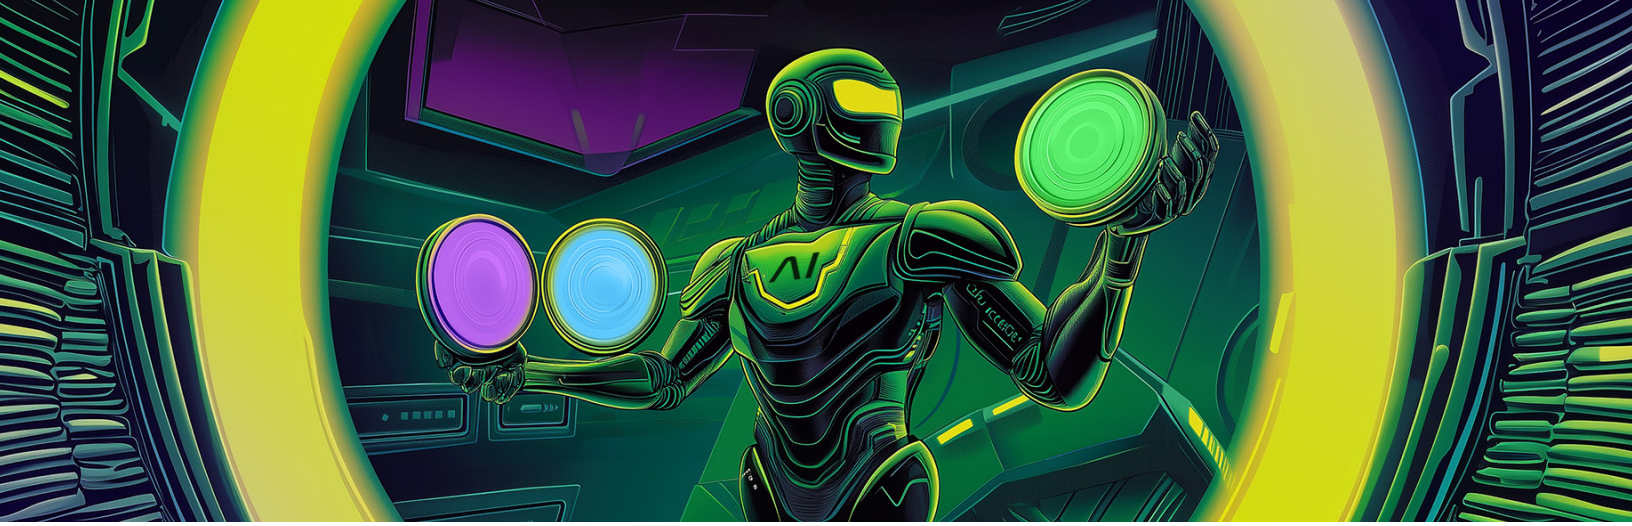

Navigating Enterprise AI: The Three Circles of AI Adoption

According to **Sonali Bhagwat**, AI adoption in enterprises can be understood through three concentric circles, each representing a different stage of integration and complexity. These circles provide a framework for understanding how companies use AI and the unique challenges they face at each level.

Circle 1 - Inner Circle: Companies Using AI Tools and Services

In the innermost circle are companies beginning their AI journey by adopting existing tools and services. These organizations typically use SaaS applications with embedded AI functionalities, such as Grammarly for writing assistance or Workday for HR recommendations. This stage represents the initial encounter with AI's potential, where enterprises explore how AI can streamline operations, improve decision-making, and enhance productivity.

For most organizations, this is where they dip their toes into AI, testing its capabilities without fully committing to building their models. However, even at this early stage, there are risks to be aware of. Introducing AI into existing processes can bring unforeseen challenges, particularly around data security and the ethical use of AI-driven insights.

Circle 2 - Middle Circle: Companies Building AI Applications for Specific Use Cases

The middle circle consists of more advanced companies in AI adoption. These organizations are not just using off-the-shelf AI tools—they are building custom AI models tailored to their business needs. This is where enterprises start to see the value of AI as they create applications that can significantly impact their operations.

At this stage, companies are developing AI applications that integrate vector databases and leverage models from providers like OpenAI and Anthropic. The focus here is on augmenting decision-making processes and enhancing business capabilities through AI-driven insights. However, with this increased customization comes a greater responsibility to manage data securely and ethically.

> "Building custom AI models tailored to specific business needs is where enterprises start to see the true value of AI, but it also requires a more nuanced approach to data security."
>
> **— Karen Lopez**

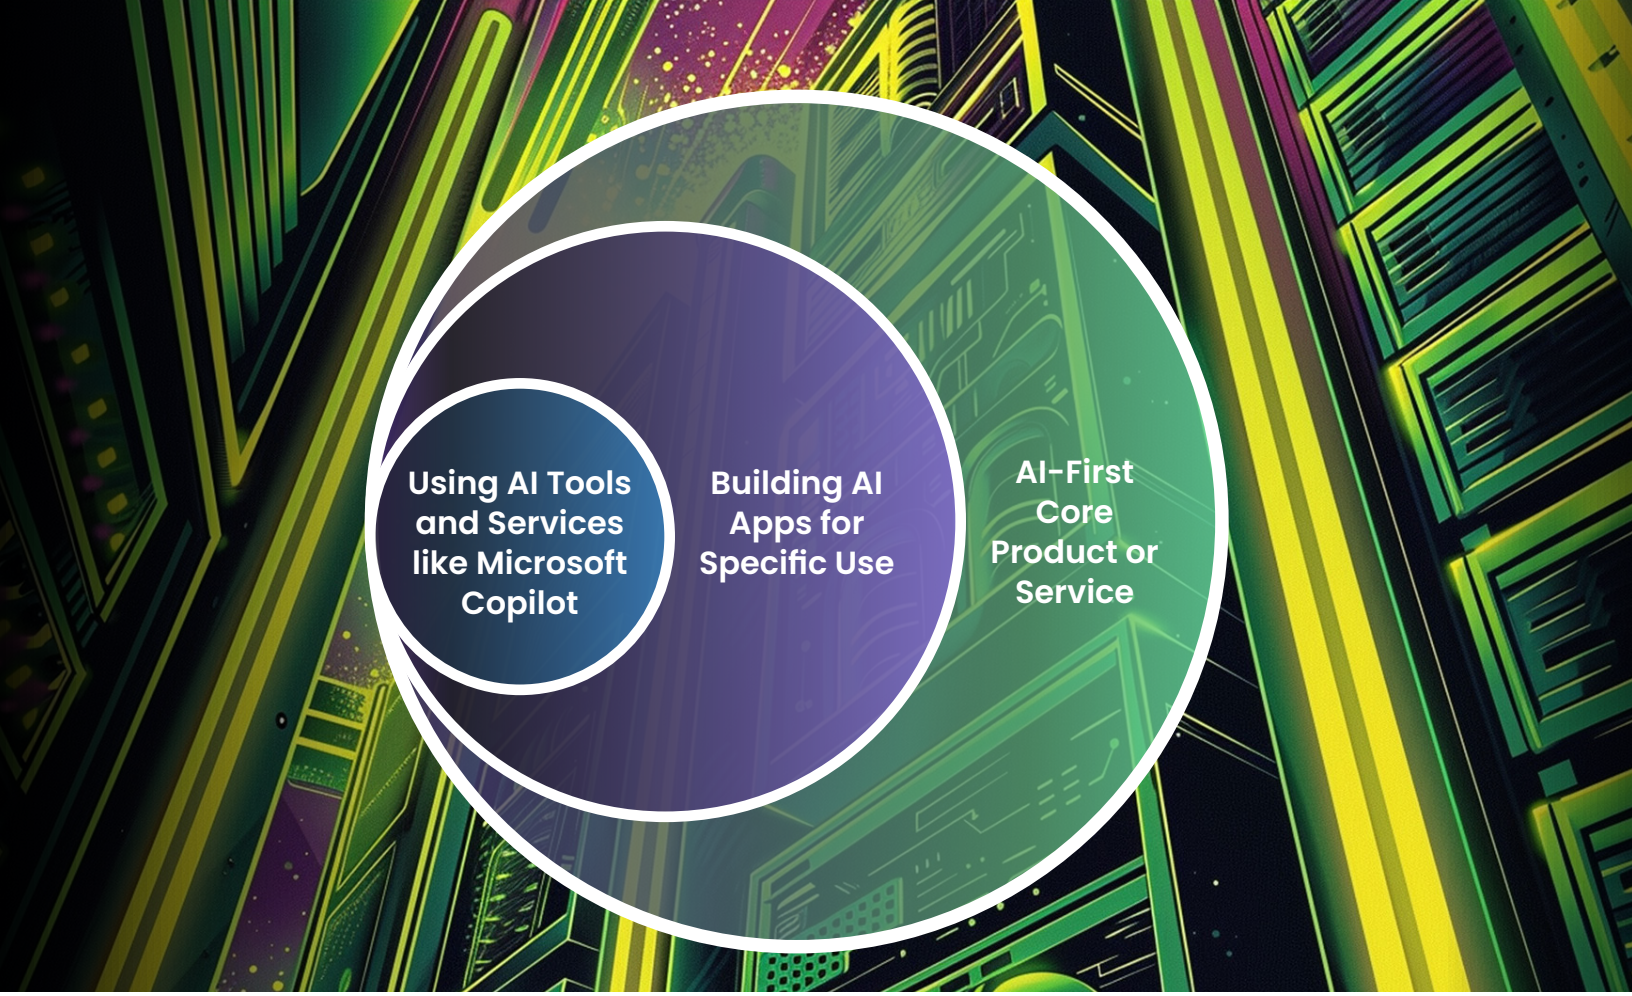

Circle 3 - Outer Circle: AI-First Companies Where AI is Integral to the Core Product or Service

The outermost circle represents companies where AI is not just a tool but the foundation of their business model. **Sonali Bhagwat** explained that these AI-first companies deeply integrate AI into their core products and services. For example, a well-known software company's AI capabilities allow users to generate images from simple text prompts, demonstrating how AI can create new customer value.

The challenges are even more pronounced in this circle. These companies must navigate complex issues around data sovereignty, ethical AI use, and the transparency of AI-driven decisions. As AI becomes integral to their offerings, ensuring the security and integrity of the data used by these AI systems is paramount.

Key AI Applications

Across these circles, the workshop participants discussed a range of AI applications that enterprises are currently using. These include custom models for specific business functions, chatbots for customer interaction, and data augmentation and manipulation tools. While these applications offer significant benefits, they also introduce new security concerns. **Arthur Hedge** noted, "***We're using AI to enhance everything from data manipulation to document generation, but with every new application, new security concerns arise."***

Understanding the landscape of AI adoption in organizations is the first step in addressing these concerns. As companies move through these concentric circles, they must continually reassess their data security strategies to keep pace with AI's evolving capabilities and risks.

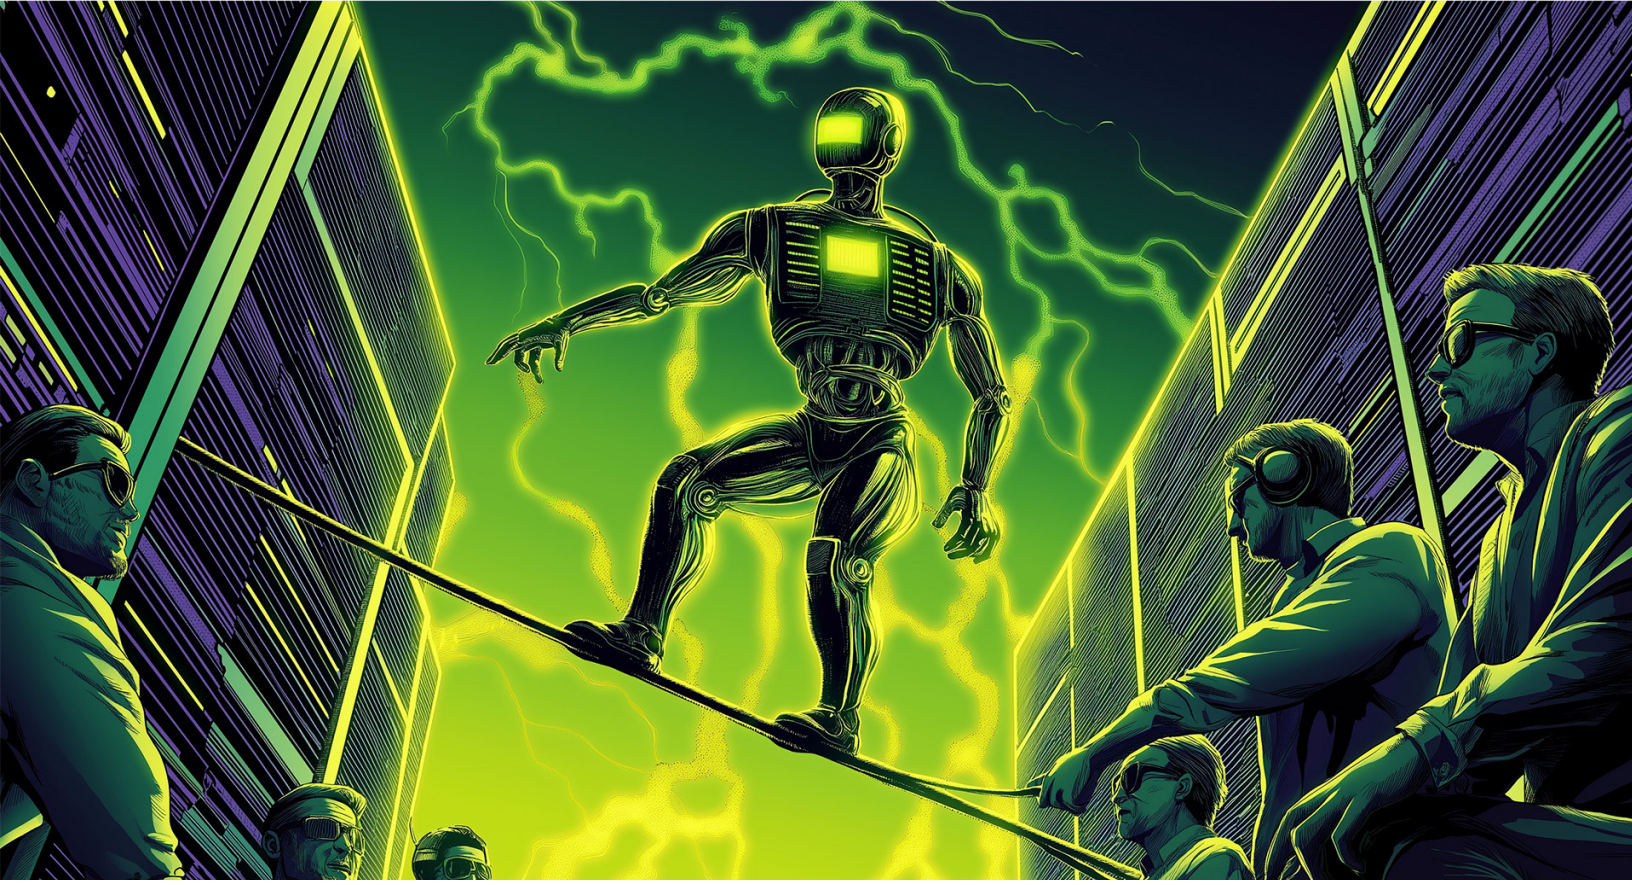

Data Security Risks in AI Applications

As enterprises increasingly integrate AI into their operations, new data security risks emerge. This section explores some of the most pressing data security threats associated with AI applications and the strategies to mitigate them.

Potential Data Security Threats

One of the most significant risks is cross-pollinating sensitive data across different clients or projects. As AI tools like Copilot gain deeper access to enterprise systems, the possibility of inadvertently mixing data from separate clients becomes a significant concern. As **Arthur Hedge** pointed out, ***"One of my biggest fears is the inadvertent cross-pollination of data between different clients, especially when using AI tools like Copilot that have deep access to our systems."*** This risk can lead to severe breaches of confidentiality and trust, making it imperative to manage how AI systems access and use data carefully.

Another critical issue is unauthorized data access due to embedded AI tools within enterprise software. As AI becomes more integrated into everyday business processes, controlling who can access sensitive data becomes increasingly complex. **Karen Lopez** highlighted this challenge: ***"The more integrated AI becomes within enterprise software, the harder it is to prevent unauthorized access. It's a constant battle to stay one step ahead."*** The dynamic nature of AI tools, which often evolve rapidly, means that traditional security measures may not be sufficient to protect against new vulnerabilities.

Additional Risks

Other significant data security threats raised during the workshop include:

Data Leakage: AI applications require access to large datasets, which increases the risk of data leakage, particularly if misconfigurations or vulnerabilities expose sensitive information. This can occur during processing and storage, making proper security protocols essential.

Insider Threats: As AI systems access more critical data, the risk of insider threats grows. Employees or contractors with access to AI tools may unintentionally or deliberately misuse them, leading to data breaches or unauthorized manipulation.

Data Poisoning: AI models are vulnerable to attacks on the data they are trained on. In some cases, attackers can introduce false or malicious data into AI training datasets, causing the AI to make incorrect or harmful decisions. This risk is particularly relevant as organizations increasingly rely on AI for decision-making.

Non-compliance with Data Residency and Sovereignty Laws: As AI systems often process data across borders, there's a heightened risk of non-compliance with data residency and sovereignty laws. Sensitive data may inadvertently be transferred to regions with weaker data protection regulations, exposing organizations to legal and regulatory risks.

Inadequate Data Retention and Deletion Policies: Many AI systems store large amounts of data for analysis, which can become a risk if proper data retention and deletion policies are not enforced. This poses a compliance issue with laws like GDPR and a security risk if old data is not correctly disposed of.

Risk Management Approaches

To mitigate these risks, organizations need to adopt sophisticated risk management strategies. One approach is implementing data classification and labeling systems. By categorizing and labeling data, enterprises can control which data AI tools can access, reducing the likelihood of unauthorized access or mishandling. **Arthur Hedge** emphasized the importance of this approach: ***"We're implementing a data classification and labeling system to ensure that our AI tools only access the data they're supposed to, which is critical to preventing security breaches."***

Additionally, organizations must focus on controlling data flow and access in AI tools like Copilot and Perplexity. While these tools offer significant value, they need to be managed carefully to prevent potential data leaks or misuse. **Jeff Farinich** noted the delicate balance required: ***"Tools like Perplexity and Copilot offer immense value, but they must be carefully managed to avoid potential data leaks or misuse."*** Effective management includes setting strict access controls, monitoring data usage, and ensuring that AI tools operate within clearly defined boundaries.

As AI advances and becomes more embedded in business environments, the risks associated with its use will evolve. Organizations can safely utilize AI while protecting their crown jewels by understanding these risks and implementing proactive risk management strategies.

Data Governance and Compliance in AI

Security and data leaders must navigate a complex landscape where traditional regulations intersect with emerging AI technologies. This section explores the key challenges and strategies for ensuring AI applications align with data protection laws and uphold transparency and data sovereignty principles.

Compliance with Data Protection Laws

While AI-specific regulations are still evolving, organizations must act quickly. Instead, they are applying existing data protection laws, such as GDPR, CPRA, and PCI-DSS, to their AI initiatives. **Jeff Farinich** encapsulated this approach: ***"While we don't have AI-specific regulations yet, we apply the same rigorous standards as we do with sensitive data to ensure compliance across the board."*** This means treating AI-generated data and AI-driven processes with the same care and scrutiny as other sensitive information, ensuring that privacy and security standards are consistently upheld.

A significant challenge in this context is the role of internal legal teams and their collaboration with technical experts. AI is rapidly evolving, and legal teams are often on a steep learning curve. **Karen Lopez** highlighted this issue: ***"A major challenge we face is the gap between our legal team's expertise and the technical nuances of AI. It's a learning curve for all of us."*** Bridging this gap requires ongoing collaboration and education, where legal and technical teams work together to develop legally sound and technically feasible policies.

Data Sovereignty and Transparency

For global enterprises, data sovereignty is a critical concern. Companies must ensure that their data practices comply with the laws of each country in which they operate, particularly regarding AI. **Sonali Bhagwat** emphasized this importance: ***"Data sovereignty is non-negotiable, especially for companies operating across borders."*** This commitment helps adhere to legal requirements and builds trust with customers and stakeholders by demonstrating respect for data privacy.

Transparency in AI decision-making is another essential aspect of governance. As AI becomes a more integral part of decision-making processes, organizations must ensure that these decisions are transparent and understandable. **Sonali Bhagwat** suggested a practical approach: ***"Every AI-driven decision should come with a 'nutritional label' indicating how much AI influenced it, ensuring full transparency."*** By providing this clarity, organizations can help users and stakeholders understand AI's role in shaping outcomes, which is crucial for maintaining trust and accountability.

Data governance and compliance require a proactive approach when dealing with AI. Organizations must comply with existing laws and anticipate future regulations while ensuring their AI practices are transparent and respect data sovereignty. By fostering collaboration between legal and technical teams and maintaining a commitment to transparency, companies can navigate the complexities of AI governance and build a foundation of trust and compliance.

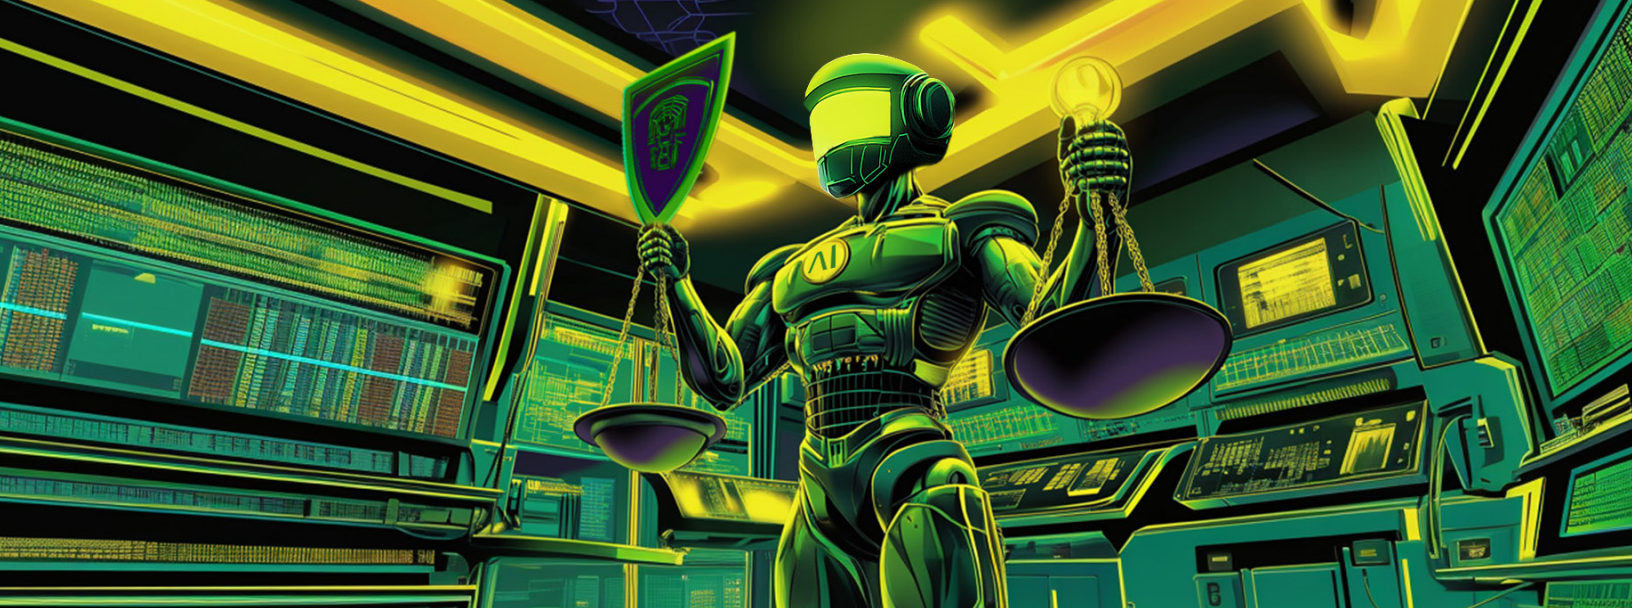

Balancing AI Innovation with Data Security

Pushing the boundaries of what AI can achieve is an exciting prospect for any organization, but it cannot come at the cost of security. As many leaders in the workshop emphasized, it's about finding that "sweet spot" where AI innovation and security can coexist. One approach is implementing AI in areas with the most immediate impact while maintaining strict controls over data access and usage. By doing so, companies can explore AI's potential without exposing themselves to unnecessary risks.

The importance of a balanced approach cannot be overstated. **Arthur Hedge** noted, ***"AI should augment, not replace, human judgment. We need to maintain a balance where AI complements human decision-making without taking over entirely."*** This perspective conveys the need for AI to be a tool that enhances human capabilities rather than diminishes human oversight's role. By positioning AI as a partner rather than a replacement, organizations can ensure that critical decisions are still guided by human insight, with AI providing the necessary support.

Case Studies and Best Practices

Workshop participants shared several examples of navigating the delicate balance between AI innovation and data security within their organizations:

Arthur Hedge discussed his organization's approach to AI-driven document generation: ***"We use AI to help draft policy documents and generate code, but we're meticulous about ensuring that the AI doesn't access sensitive client information. Everything is closely monitored."*** This method emphasizes strict oversight and clear boundaries on what AI can and cannot access, prioritizing data protection at every step.

Karen Lopez shared how her team is testing AI in data analysis while being cautious about data sovereignty: ***"We're using AI to analyze data sets for anomalies, but we keep everything within the country of origin to respect data sovereignty laws. It's a step-by-step approach, and we're learning as we go."*** Her case highlights the importance of aligning AI innovation with compliance requirements, particularly in global operations.

These case studies highlight the importance of starting small, learning from each implementation, and gradually expanding AI's role as confidence in its security grows. By taking a measured approach, organizations can build a foundation of trust in their AI systems, ensuring that innovation does not outpace security.

Moving Forward with AI in Data Security and Governance

Looking forward, it's clear that AI will become increasingly central to how enterprises function. **Sonali Bhagwat** noted, ***"AI is here to stay, and its role in enterprise operations will only grow. But with that growth comes increased responsibility to manage it securely."*** While AI will drive efficiencies and unlock new capabilities, it will also demand a more sophisticated data security and governance approach.

The path forward for enterprises embracing AI should be deliberate and well-considered. **Karen Lopez** emphasized the importance of a cautious, structured approach: ***"My recommendation? Start small, build a strong governance framework, and scale responsibly. AI isn't a race—it's a journey."*** This advice underscores the need for a foundation of robust governance policies before expanding AI initiatives. Organizations can learn and adapt by starting with manageable projects, ensuring their AI deployments are secure and compliant.

The Evolution of AI Regulations:

Jeff Farinich mentioned, ***"As AI continues to evolve, we can expect more regulations specifically targeting AI. Preparing for this now, by aligning our practices with the most stringent data protection laws, will make future compliance much easier."*** This highlights the need for proactive measures, anticipating stricter regulations, and preparing accordingly.

AI and Ethical Considerations:

Karen Lopez touched on the ethical dimensions of AI, saying, ***"Beyond just security, we need to think about the ethical implications of AI in decision-making. It's not just about protecting data—it's about using AI in ways that are fair and just."*** This forward-thinking perspective emphasizes incorporating ethical guidelines into AI governance frameworks.

Integrating AI with Existing Systems:

Arthur Hedge discussed the importance of integration, ***"The future of AI in enterprises isn't just about standalone AI applications—it's about how well AI integrates with our existing systems. Seamless integration will be key to maximizing AI's value while maintaining control over data security."*** This insight underscores the importance of ensuring that AI works harmoniously with existing technologies, which will be crucial for secure and effective AI adoption.

The future of AI in data security and governance is not just about technology—it's about strategy, planning, and ongoing vigilance. Organizations that approach AI with a clear focus on security and governance will be better positioned to reap the benefits of AI while minimizing its risks. Organizations can ensure their AI journey is innovative and secure by fostering a culture of responsibility and careful growth.

Charting the AI Path

As we've explored throughout this white paper, the key to successfully leveraging AI in the enterprise is balance—balancing innovation with security, transparency with functionality, and risk with reward. AI offers tremendous potential to transform business operations, but with that potential comes the responsibility to protect data, ensure compliance, and maintain ethical standards.

The insights shared by industry leaders during our workshop underline the importance of a thoughtful, strategic approach to AI adoption. Whether it's implementing data governance frameworks, carefully managing AI's integration into existing systems, or anticipating future regulatory requirements, the path forward requires caution and creativity.

But this is just the beginning of the conversation. AI's challenges and opportunities are evolving rapidly, and staying ahead will require ongoing dialogue and collaboration. We invite you to join us in our upcoming webinar, where we'll dive deeper into these topics, share new developments, and continue to explore how we can responsibly harness the power of AI.

Your participation and insights are crucial as we navigate this journey together. We must ensure that AI drives innovation while safeguarding the integrity and security of our enterprises. Let's keep the conversation going.

Interested in the latest in data security and governance?

For insights from the Data Security Creative Council Workshop, details about our founding members, or to join our dynamic community of data and security professionals, click here. Don't miss out on our upcoming content, webinars, workshops, and events designed to keep you at the forefront of data security innovation.