AI COMPLIANCE AND GOVERNANCE

The proliferation of AI in the workplace is placing unprecedented demands on compliance teams to understand a new technology, its risks and the controls it needs. This session will look at what organizations are doing and provide insights into how to improve your AI governance efforts.

Robert Hughes, Chief Information Security Officer, RSA

Justin Lofquist, Director of CMEP Core Services, NERC

SAI360

The Evolution of AI

Brief History of AI, and what can we expect in the future

- **1950s–1970s:** Early rule-based AI - Turing poses 'Can machines think?'
- **1980s–1990s:** AI winter hits, machine learning emerges as a discipline
- **2000s–2010s:** Big data, GPU computing, and deep learning reignite AI
- **2017:** Transformer architecture unlocks language models (LLMs) at scale
- **2022–Present:** Generative AI reaches mainstream; enterprise adoption accelerates globally

AI Risk Factors

The fundamentals can guide us

- Data security – protect from loss / Breach
- Supply chain / Third Party Risk Management
- Managing Shadow IT
- Identity and Role-based Access
- User awareness, training and culture
- IT Tool Lifecycle
- SDLC – Peer review and secure code practices

AI Risk Factors

What feels new

- Compressed technology adoption cycle - Faster
- Third party integrations and data flows - Attack surface expands
- It may not all be logged and traceable
- Lot of non-human Identities (NHI)
- New regulatory frameworks
- Pressure for more automation / autonomy

Show me the logs

Deterministic vs. Probabilistic

- Deterministic systems: same input always produces the same output — auditable and predictable
- LLMs are probabilistic: outputs vary by design, making reproducibility a compliance challenge
- Traditional controls (input/output validation, unit tests) insufficient for probabilistic AI
- Auditors must shift to statistical sampling, red-teaming, and behavioral benchmarks
- Governance must account for 'hallucinations,' model drift, and prompt-sensitivity over time

Is there peer review?

The Human in the Loop

- Humans cannot review every AI decision at scale
- Defining and validating the guardrails: thresholds, policies, and constraints
- Exception-based monitoring
- Document the oversight model itself — auditors need evidence of the governance framework

Is AI doing what it is supposed to?

Measuring the system

- Intent – what was it? Documentation
- Duty of Care, Duty of Loyalty
- Who is accountable for AI decisions?

Approaching the Singularity

Vulnerability management disruption – Claude Mythos

- It may not be tomorrow, but it's coming
- An AI capable enough to find vulnerabilities faster / better than humans and existing scanners
- Mythos is not alone – ChatGPT 5.5 is coming – the others will catch up
- Expectation: Good security practices win out:
  - Defense in Depth, Zero Trust
  - Good patching practices – be ready for more patches, faster!

Approaching the Singularity

Other new threats for defenders

- Social engineering is better and easier to do
- Adaptive malware is (more) trivial to create – signatures less useful
- Chaining exploits in non-human-intuitive ways
- The noise, it will be louder

Approaching the Singularity

Help for defenders

- Natural language data queries and reporting – threat hunting
- Easier to manage and coordinate data
- Anomaly detection
- Compliance documentation
- Compliance research and control framework management

Securing AI Agents

Emergent Behavior

Cascading Failures

Obscured Casual Providence

Decision Fatigue

- Governance, risk, and compliance (GRC)
- Identity and access management (IAM)
- Data security and privacy
- Application security
- Threat management
- Zero Trust architecture

AI Risk Control Frameworks

- NIST AI RMF
- EU AI Act
- OWASP Top 10 for LLMs
- MITRE ATLAS
- ISO/IEC 42001

- Anchor strategy in NIST and ISO
- Use EU and OWASP to deepen compliance
- Start with high-density controls for broad coverage

AI Governance

Practices that business leaders adopt to incorporate purpose, culture, action, and assessment to ensure AI delivers desired business outcomes, is responsibly used, and complies with applicable regulations.

Product Mindset For Data and AI

- Treat data and AI assets as products.
- Build your capabilities with a living roadmap
- Assign product roles to drive the product

Governance Versus Democratization

Governance

- Align by design to mitigate misalignment risk
- Build AI governance on four pillars: purpose, culture, action, and assessment
- Establish a cross-functional governance team to create a governance framework
- Embed disciplined audits and automated controls in every data and AI pipeline

Democratization

- Launch an enterprise data and AI product catalog
- Deploy low-code, no-code, and conversational (NLQ/genAI) analytics environments
- Run a company-wide data literacy program

Governance: Defense versus Offense

Defense

- Satisfy existing and impending regulatory requirements
- Adhere to responsible AI principles
- Mitigate attacks on AI itself

Offense

- Improve customer trust.
- Preserve corporate values
- Drive revenue growth

Governance to Drive Strategic Outcomes

- Security
  - Prevent security breaches through frameworks like system monitoring, policy enforcement technologies, and cross-functional oversight, thus preventing significant financial losses and reputational damage.
- Privacy
  - Align governance policies with frameworks like GDPR, HIPAA, and CCPA to protect individual privacy rights, manage data acquisition, and ensure appropriate access, building customer and employee trust.
- Compliance
  - Approach evolving regulatory requirements with a future-focused mindset to transform compliance from a defensive barrier into an offensive business strategy that enables innovation within structured requirements.

Governance to Drive Strategic Outcomes

- Self-Service
  - Break down silos and accelerate data-driven decision-making by empowering employees across all roles to independently access, discover, and use high-quality data and AI products regardless of their technical expertise.
- Discovery
  - Enable enterprise-wide visibility, trust, and understanding of data and AI assets through comprehensive catalogs, metadata management, and searchable platforms.

Four Pillars for AI Governance Practice

- Purpose:
  - Specify business outcomes that the AI system should achieve
- Culture:
  - Embed AI governance into existing practices and roles.
- Action:
  - Enable and automate AI stewardship tasks and processes
- Assessment:
  - Observe, measure, and communicate AI impact.

Advance AI Governance

- Get started by going beyond technical data science proficiency.
- Establish a baseline, optimize, and retire AI you build, buy, and compose.
- Get executives involved early.
- Federate AI governance within lines of business.

AI Governance Solutions

- Market Definition
- Business Value
- Market Maturity
- Market Dynamics
- Notable Vendors
- Where the market is going next
- Drive revenue growth

Use Cases and Q&A / Discussion

- **RSA**
  - Background
  - Unique challenges as a services provider (supply chain)
- **NERC**
  - Background
  - Unique challenges as a regulator

Thank You!