



Auditing the Algorithm

Ethics and Compliance in the Age of AI

PRESENTER

Julie DeMuth Mellendorf, AIGP, CIPP/US

Crowe Advisory LLC

National Council of Philippine American
Canadian Accountants (NCPACA) Convention

September 4, 2026

Educational content; not legal or regulatory advice.

Disclaimer

- This presentation is provided for general informational and educational purposes only and does not constitute legal, tax, accounting, or other professional advice.
- It should not be relied upon as a substitute for consultation with qualified legal, compliance, or risk advisors familiar with your organization's specific facts and circumstances.
- Nothing in this material creates an attorney-client, advisory, or other professional relationship between Crowe LLP and Crowe Advisory LLC and any recipient.
- Views expressed are those of the presenter and do not necessarily reflect the official position of Crowe LLP and Crowe Advisory LLC.

Meet Today's Speaker



Julie DeMuth Mellendorf, AIGP, CIPP/US

Quality and Risk Management Leader, Crowe Studio

AI Ethics and Governance

Chair, Crowe AI Risk and Ethics Subcommittee

Julie.Mellendorf@crowe.com



AGENDA

What We Will Cover

A practical guide to evaluating and governing AI

ONE

Foundation

Why disciplined oversight matters more than assurance alone

TWO

Risk and Ethics

Purpose, data, reliability, confidentiality, and human responsibility

THREE

Governance in Practice

Intake, accountability, evidence, and monitoring

FOUR

High-Level Audit Lens

Implications when AI touches reporting or assurance

AI Is Already in the Workflow

The practical question is whether oversight, evidence, and accountability will keep pace with adoption.

WHERE IT IS ALREADY SHOWING UP



Everyday productivity

Drafting, research, summaries



Client service

Analysis, recommendations, deliverables



Business operations

Finance, HR, legal, marketing



Embedded vendor AI

Features inside existing platforms



High-impact decisions

Scoring, eligibility, access, approvals

A REPEATABLE PATH FROM DISCOVERY TO DEFENSIBLE OVERSIGHT

DISCOVER

Find actual use, pilots, and embedded features

INVENTORY

Record purpose, owner, users, data, and vendor

ASSESS

Evaluate impact, behavior, data, and dependencies

APPROVE

Set conditions, controls, and decision authority

MONITOR

Track performance, exceptions, and change

IMPROVE

Re-review, remediate, refine, or retire

Governance is the thread through every step, ownership, proportionate controls, traceable evidence, and escalation.

AI can accelerate the work; it cannot own the decision.

2026: AI governance has moved from principles to proof

EU transparency rules are now being enforced.

U.S. states are imposing documentation and impact-assessment duties.

Existing privacy and securities laws are being applied to AI.

Fraud standards are adapting to synthetic evidence and impersonation.


Regulatory requirements continue to change after deployment.

FOUNDATION

Building a Defensible AI Program


Begins with the questions leadership should be able to answer.

A Defensible Position Starts With Ten Questions



Four Building Blocks Create a Minimum Viable Program

Start with a usable system, then mature it as risk and adoption grow.



Direction → boundaries → decisions → evidence → learning

TEXAS EXAMPLE: Its Responsible AI Governance Act (effective Jan. 1, 2026) lets the Attorney General demand this same documentation, purpose, training data, outputs, performance, limitations, monitoring, making it evidence of reasonable care.

Integrate AI Oversight Into Existing Functions

Coordinate the work already performed by risk, legal, privacy, security, data, vendor, and assurance teams



Lean On Established Frameworks

Regulators and examiners are looking for risk management application

RISK FUNCTIONS



NIST AI RMF

- Voluntary and use-case agnostic
- Govern • Map • Measure • Manage
- Useful backbone for risk conversations

Best use: common risk language

MANAGEMENT SYSTEM



ISO/IEC 42001

- Policies, objectives, roles, and processes
- Plan • Do • Check • Act
- Useful for enterprise-wide operating discipline

Best use: operating discipline

FINANCIAL SERVICES



FS AI RMF

- NIST-aligned and sector-specific
- Adoption-stage questionnaire + control matrix
- 230 control objectives and evidence examples

Best use: sector-tailored controls

AUDIT & ASSURANCE



ISACA AI Audit Toolkit

- Built for reviewing AI governance
- Governance • Risk • Control testing
- Maps to COBIT and existing IT audit work

Best use: independent assurance

Choose a backbone → map it to existing processes → borrow only what improves decisions and evidence

Unapproved AI Use Creates Multiple Risks at Once



Confidentiality & Privacy

Confidentiality and privacy can be compromised at input.



Output Accuracy

Outputs can be inaccurate, biased, or fabricated.



IP & Contractual Risk

Prompts and outputs may create intellectual-property or contractual issues.



Human Accountability

Human owners remain accountable for decisions and records.

Several risks may arise at once—but data can leave control before an output is generated.

NORTH CAROLINA EXAMPLE: Its Responsible Use of AI Framework requires agencies to inventory AI tools and assess risk before and after deployment — work that is impossible for a tool nobody knows is in use, leaving unknown data exposure, no owner, no assessment, and no audit trail.

RISK & OVERSIGHT

What the Review Should Cover

A structured review of purpose, data, behavior, controls, outcomes, and accountability.

Audit the System, Not Just the Output

A disciplined review follows purpose, data, behavior, change, and accountability.


DIMENSION	BUSINESS QUESTION	GOVERNANCE EVIDENCE	REVIEW TEST
Purpose / decision What role should AI play?	- Owner, intended outcome, users, decision authority - Recorded scope and risk tier	- Is AI suitable for this task and impact? - Are human responsibilities clear?	APPROPRIATENESS Confirm purpose, authority, and boundaries.
Data / permissions What data enters, leaves, or persists?	- Source, authority, minimization, retention, vendor terms - Lineage and access record	- Is data use permitted and proportionate? - Can the data path be reconstructed?	AUTHORIZATION Verify permissions, provenance, and handling.
Behavior / reliability How can the system fail, bias, or mislead?	- Testing, limitations, human review, exception handling - Baseline and acceptance criteria	- Are outputs reliable enough for the use? - Are contrary signals investigated?	CHALLENGE Test performance, fairness, and failure modes.
Change / accountability What can change after approval?	- Version, vendor, data, incident, and monitoring history - Named re-review triggers	- Can another reviewer follow the decision? - Can issues be escalated and corrected?	TRACEABILITY Retain the path from purpose to decision and monitoring.

When AI touches financial reporting or an audit, professional standards still govern reliance, evidence, documentation, and judgment.

CANADA MODEL: Its mandatory Algorithmic Impact Assessment evaluates the system, decision, affected population, and data, then assigns one of four impact levels — required safeguards scale with the level, and it is revisited when scope changes: a model for proportionate, risk-tiered review.

A High-Level Review Connects Purpose to Evidence

The goal is defensible oversight—not a technical code examination.



Governance evidence retained with the use case

Purpose and owner

Data and vendor review

Testing and validation

Human oversight

Escalation record

Monitoring evidence

If AI affects reporting or assurance, apply the relevant professional standards to evidence, documentation, and judgment.

DELOITTE EXAMPLE: Fabricated citations and a fake court quote forced a correction that disclosed AI-tool use — verify every citation with a named reviewer.

Synthetic Content Raises Evidence and Fraud Risk

Deepfakes and generated records raise authenticity risk across decisions, investigations, reporting, and assurance.

EMERGING SIGNAL — AND THE GOVERNANCE RESPONSE

Deepfake instructions or approvals

Voice or video may imitate an authorized executive.

Response: independent verification + corroboration

Synthetic documents and records

Polished documents may lack authentic provenance.

Response: source, metadata, and counterparty checks

Compromised communications

Spoofed domains or channels can redirect decisions.

Response: control the channel and use alternate verification

Fabricated citations and analysis

AI can invent authorities, sources, and calculations.

Response: verify primary sources and recalculate

Escalate when:

identity is uncertain • provenance breaks • records conflict • access or instructions change

RESPONSE BY SIGNAL

ROUTINE

Corroborate source, authority, content, and metadata.

ELEVATED

Use alternate verification and a second reviewer.

CRITICAL

Stop, escalate, preserve records, and reassess fraud or misconduct risk.

Unexpected ease, polish, or consistency is a reason to verify—not proof of authenticity.

FBI / FINCEN: In May 2025 the FBI warned of AI-generated voice messages impersonating senior officials, and FinCEN has flagged rising deepfake-identity fraud — the control response should be process, not voice recognition: independent callback, verified contacts, dual approval, MFA, and escalation for unusual transactions.

Not Every AI Use Case Deserves the Same Review

Risk and strategic value should determine the pathway, controls, and approval level.

01



Internal productivity

Drafting, search, summarization, coding, or analysis?

02



External interaction

Client, customer, employee, or public-facing output?

03



Consequential decision

Does AI influence eligibility, access, pricing, rights, or obligations?

04



Sensitive data

Does the use involve confidential, personal, privileged, or regulated data?

05



Third-party dependency

What do vendor terms, transparency, security, and change controls allow?

06




Autonomy & change

Can the system act, adapt, or change without timely human intervention?

Higher impact, sensitive data, limited reversibility, and autonomous action justify stronger review and evidence.

Prioritize With Two Lenses: Strategic Value and Risk

The same intake should inform both portfolio priority and control intensity.



Examples are illustrative; classification depends on institutional context.

Use the matrix to decide:

- What moves first
- How much review is justified
- Where to escalate or stop

Reassess when the use, data, vendor, or model changes.

GOVERNANCE & READINESS




Turning Review Into Routine

How owners, policies, evidence, and monitoring work together.

Clear Roles Create the Right Level of Confidence

Separate ownership, challenge, independent review, and formal assurance.



Confidence depends on clear roles, scope, independence, criteria, and evidence.

AI Oversight Works in Layers


Each layer answers a different question and leaves evidence for the next.

LAYER	PRIMARY OWNERS	CORE EVIDENCE	PRACTICAL QUESTION
Enterprise governance Direction, boundaries, and accountability	• Leadership, board, risk committee, and legal • Strategy, risk appetite, policy, and reporting	• Framework, policy, inventory, and committee records • No new data sharing; vendor assessed	ASK Are priorities, boundaries, and accountability clear?
Use-case governance Purpose, impact, and decision rights	• Business owner, process lead, and risk reviewers • Assessment, approval conditions, testing, and decision record	• Use-case record, risk tier, test results, and approvals • False positives create customer harm	ASK Is the pathway proportionate to impact and risk?
Data & vendor governance Inputs, providers, access, and change	• Data, privacy, security, procurement, and vendor management • Due diligence, authority, access, retention, and change controls	• Data lineage, vendor record, access log, and change history • Adverse-action reasoning required	ASK Can the organization trace what enters, leaves, and changes?
Independent oversight Challenge, monitoring, and assurance	• Compliance, risk, internal audit, and external assurance when appropriate • Monitoring, exceptions, incidents, remediation, and review results	• Independent review results and management responses • Vendor terms and retention unclear	ASK Can a reviewer follow the evidence and challenge the conclusion?

Connect the layers: enterprise direction sets boundaries; use-case evidence shows how decisions operate in practice.

A 90-Day Plan Can Turn AI Oversight Into Routine

Build usable guardrails, decisions, evidence, and monitoring in three practical sprints.



Start with the work already happening; govern it, evidence it, monitor it, and improve it.

Five Practices Make AI Oversight Defensible

The objective is not perfect control; it is consistent, risk-based, explainable decisions.

FIVE PRACTICES



Visibility comes first

Inventory actual use, owners, data, and vendors.



Risk drives effort

Apply stronger review where impact is higher.



Humans own decisions

Preserve authority, challenge, and override.



Evidence makes review defensible

Record purpose, assessment, approval, and testing.



Monitoring closes the loop

Track change, exceptions, incidents, and remediation.

A REPEATABLE DISCIPLINE FOR AI OVERSIGHT

DISCOVER

Find actual use

CLASSIFY

Purpose, impact, data

CHALLENGE

Assumptions and controls

EVIDENCE

Decision and rationale

DECIDE

Human accountability

MONITOR

Change and outcomes

Visibility • proportionality • accountability • evidence • monitoring

“Auditing the algorithm” is a management discipline: ask, challenge, evidence, decide, and revisit.



Thank You.

Questions?



Julie DeMuth Mellendorf, AIGP, CIPP/US

Quality and Risk Management Leader, Crowe Studio

Julie.Mellendorf@crowe.com

crowe.com

"Crowe" is the brand name under which Crowe LLP, Crowe Advisory LLC, and their respective subsidiary entities provide professional services. Crowe LLP and Crowe Advisory LLC (and their respective subsidiary entities) practice in an alternative practice structure in accordance with the AICPA Code of Professional Conduct and applicable law, regulations, and professional standards.

Crowe LLP is a licensed independent CPA firm that provides attest services to its clients. Crowe Advisory LLC and its subsidiary entities provide tax, advisory, consulting, and other nonattest services to their clients. Crowe Advisory LLC and its subsidiary entities are not licensed CPA firms. Crowe LLP and Crowe Advisory LLC are independent members of the Crowe Global network. Crowe Global and its member firms, including Crowe LLP and Crowe Advisory LLC, are not a worldwide partnership. Crowe Global and each member firm are separate legal entities. Services are delivered by the member firms; Crowe Global is a Swiss organization and does not render professional services to clients and does not have an ownership interest in Crowe LLP, Crowe Advisory LLC, or any other member firm. Crowe Global and its member firms are not agents of, and do not obligate, one another and are not liable for one another's acts or omissions.

The information contained in this document is not intended to constitute audit, tax, accounting, advisory, consulting, business, financial, investment, legal, or other professional advice and should not be relied upon as such. Some services may not be available to attest clients. The information is general in nature, based on authorities that are subject to change, and is not a substitute for professional advice or services. Before making any decision or taking any action that may affect your business, you should consult a qualified professional adviser. Crowe is not responsible for any loss incurred by any person who relies on the information contained in this document.

For additional information about Crowe LLP, Crowe Advisory LLC, their subsidiary entities, and Crowe Global, please visit <https://www.crowe.com/disclosure>. © 2026 Crowe Advisory LLC. All rights reserved. The Crowe name and logo are trademarks used under license from Crowe Global.