1. Consider Adoption of its Spring National Meeting Minutes

Attachment One

Commissioner Michael Yaworsky (FL)

NAIC NATIONAL ASSOCIATION OF INSURANCE COMMISSIONERS

Draft: 5/5/26

Innovation, Cybersecurity, and Technology (H) Committee
San Diego, California
March 25, 2026

The Innovation, Cybersecurity, and Technology (H) Committee met in San Diego, CA, March 25, 2026. The following Committee members participated: Michael Yaworsky, Chair (FL); Karima M. Woods, Vice-Chair represented by Dana Sheppard (DC); Heather Carpenter (AK); Mark Fowler (AL); Jimmy Harris (AR); Michael Conway represented by Jason Lapham (CO); Doug Ommen (IA); Jon Godfread represented by Colton Schulz (ND); Eric Dunning (NE); Michael Humphreys (PA); Bill Huddelston (TN); and Kaj Samsom and Mary Block (VT). Also participating were: Lori Dreaver Munn (AZ); Josh Hershman, George Bradner, and Kristin Fabian (CT); Sandra Darby (ME); Christian Citarella (NH); Tom Botsko (OH); Elizabeth Kelleher Dwyer and Matthew Gendron (RI); Amanda Crawford (TX); Scott A. White and Michael Peterson (VA); and Nathan Houdek (WI).

1. Adopted its 2025 Fall National Meeting Minutes

Schulz made a motion, seconded by Fowler, to adopt the Committee's Dec. 11, 2025, minutes (*see NAIC Proceedings – Fall 2025, Innovation, Cybersecurity, and Technology (H) Committee*). The motion passed unanimously.

2. Adopted the Reports of its Working Groups and Subgroup

   A. Big Data and Artificial Intelligence (H) Working Group

Commissioner Houdek said that during the Working Group's meeting on March 24, the Working Group adopted its Feb. 17 meeting minutes.

The Working Group next received an update on the artificial intelligence (AI) systems evaluation tool pilot, which officially started earlier in March. Commissioner Houdek said that pilot states are using the tool to support a mix of market conduct exams, financial exams, and financial analysis, and that it is part of a more general regulatory inquiry. Pilot state regulators are maintaining regular communication and coordinating with companies to include the pilot, and the Working Group will provide public updates throughout the process. The pilot summary document can be found on the Working Group's web page.

The Working Group then received a presentation on operationalizing the NAIC's *Model Bulletin on the Use of Artificial Intelligence Systems by Insurers*. The presentation explained possible approaches to implement the model bulletin through governance documentation and oversight practices.

Finally, the Working Group heard a panel discussion on AI governance trends from Scott Kosnoff of Faegre Drinker and Anthony Habayeb at Monitaur. The panel discussed emerging best practices for AI governance, highlighted the importance of cross-functional governance committees, emphasized the need for strong AI inventory management, especially for high-risk systems, and underscored the value of risk-based triage, focusing governance resources on materially risky or high-impact AI use cases. The panelists framed AI governance as an evolving risk management discipline that should be integrated into existing enterprise controls and decision-making structures.

B. Privacy Protections (H) Working Group

Director Dwyer said that the Privacy Protections Drafting Group met Nov. 7, 2025, and Dec. 3, 2025, in regulator to regulator session to consider revisions to Article 6 of the *Privacy of Consumer Financial and Health Information Regulation* (#672). Following those discussions, regulators published proposed changes as part of a longer-term effort to reassemble the model and re-expose it for public comment.

There are now three articles left to look at. Article 7 was exposed for a 30-day public comment period earlier in March. After reviewing the comments on Article 7, the work will then transition to Article 8 and Article 1, which contain the definitions that the regulators will later have to consider.

Director Dwyer further said that the Working Group did not meet at the Spring National Meeting but will continue to meet virtually. She said that later, the drafting group will submit the full revised draft of Model #672 to the Working Group, after which, the Working Group will be able to take up additional comments at the end of this year.

C. Third-Party Data and Models (H) Working Group

Lapham said the Third-Party Data and Models (H) Working Group met March 23. During this meeting, the Working Group adopted its minutes and discussed potential revisions to the third-party regulatory framework. The exposed draft framework focused on company operations with direct consumer impact, including pricing, underwriting, utilization reviews, claims, fraud, and marketing. He said that during the meeting, the Working Group reached consensus to narrow the focus to pricing and underwriting as a first step.

Lapham said this process would continue to be iterative and noted that the registration process is better described as a registry. He said that the Working Group had reached consensus on the third-party data and models registry, focusing on the governance review of the datasets and models. The registry will provide a market view of third-party activities and gather information from third parties about their data and model governance. The registry would contain information so that third parties could validate those third parties operating in the insurance space. The registry concept represents an initial step that enables regulators to verify consistent national governance standards across third parties, thereby strengthening consumer protection.

The Working Group also met March 19 and Feb. 5 in regulator-to-regulator session, pursuant to paragraph 3 (specific companies, entities, or individuals) of the NAIC Policy Statement on Open Meetings, to continue work on its goals.

D. SupTech/GovTech (H) Subgroup

Munn said that since the SupTech/GovTech (H) Subgroup focuses on regulatory tools and education, and receives feedback from companies on proprietary technology, the meetings are conducted in regulator-to-regulator session.

The Subgroup continues to use the results of the regulator survey conducted last year to identify topics that drive efficiency, reduce administrative burden, and strengthen regulatory effectiveness.

Munn said that since the 2025 Fall National Meeting, the Subgroup met March 3 to receive presentations from three insurance departments on their work to create and foster their data and analytics teams. Departments also shared their experiences interacting with insurance companies on their use of AI.

E. Data Call Study Group

Schulz said that the Study Group continues to modify its approach to meet evolving priorities in other committees. The Study Group began with a plan to develop a manual inventory of all NAIC data elements, including definitions, but quickly recognized this was a case of boiling the ocean. The Study Group's efforts are now focused on market regulation data elements, but it has an automated approach for creating the data inventory. The Study Group's work relies on metadata within the NAIC's enterprise data platform, which is an ongoing project for NAIC committee support to flesh out, including adding business descriptions for the database and column names. An initial inventory has now been generated and provisioned in NAIC's ThoughtSpot tool. It includes Market Conduct Annual Statement (MCAS) data, complaints data, and the homeowners data call. Schulz said he would be reviewing the draft inventory before providing it to regulators. He will then work with committee support to collect information about state ad-hoc data calls, including those using the NAIC's Regulatory Data Capture tool and those performed outside of the NAIC.

Schulz also said that later, committee support will incorporate financial statement data elements into the data inventory. All of this work will then be used to identify data gaps.

Commissioner Ommen made a motion, seconded by Director Carpenter, to adopt the reports of the: Big Data and Artificial Intelligence (H) Working Group (Attachment One), including its Feb. 17 and Feb. 9 minutes; Third-Party Data and Models (H) Working Group (Attachment Two), including its Feb. 26 minutes; SupTech/GovTech (H) Working Group; and Data Call Study Group. The motion passed unanimously.

3. Received an Update on the Cybersecurity Event Notification Portal and Adopted the Report of the Cybersecurity (H) Working Group

Commissioner Yaworsky discussed the ongoing work by the Cybersecurity (H) Working Group. He said that the Working Group has been leading discussions to develop a cybersecurity event notification portal. Departments of insurance (DOIs) currently receive cybersecurity event notices directly, and this application would allow the DOIs to receive such notices through the NAIC. Recognizing that some members of the Committee were relatively new to the discussion, Commissioner Yaworsky said he wanted to discuss this project to raise awareness of its components, anticipating a future Committee meeting to consider adoption.

Peterson also provided an update on the Working Group's activities. He said that the Working Group met at the 2025 Fall National Meeting to hear comments from members, interested state insurance regulators, and interested parties of the Working Group on the cybersecurity event notification portal project intake form and to hear a brief presentation on the NAIC's 2025 Cyber Insurance Market Report.

The Working Group then met Feb. 6 in regulator-to-regulator session, pursuant to paragraph 4 (internal or administrative matters of the NAIC) of the NAIC Policy Statement on Open Meetings, to discuss the revised version of the cybersecurity event notification portal project intake form and expose it for a public comment period.

The Working Group then met March 13. During this meeting, it reviewed and adopted revisions to the cybersecurity event notification portal project document following the most recent public exposure. He said that

key updates include the addition of a draft standard reporting form, clarifications that licensees will not be charged to use the portal, and refinements to the Service Organization Control (SOC) 3 related language, which is a security report the NAIC will make available to stakeholders once the portal is operational. Industry stakeholders expressed general support for the project and appreciation for collaboration while identifying issues such as regulator access, data governance, concentration, risk, and confidentiality that will be addressed as technical details are developed. He said the Working Group unanimously adopted the revised project document as exposed.

Commissioner Yaworsky stated that the Committee plans to consider its own adoption of the project proposal at an April meeting to be announced and asked that Committee members reach out to speak to Peterson if they have questions about the project.

Schulz made a motion, seconded by Director Dunning, to adopt the report of the Cybersecurity (H) Working Group (Attachment Three), including its March 13 minutes.

4. Heard a Presentation from PwC on Insurance AI Trends, Including Agentic AI Applications

The Committee heard a presentation from PwC on insurance industry AI trends. Scott Froseth (PwC) described how many insurers remain in a transitional phase, constrained by legacy technology architectures and data readiness issues. As a result, most organizations are currently deploying AI through incremental, point-solution use cases rather than undertaking broader operating model transformation. He noted that while leading insurers are beginning to experiment with more sophisticated agent-based workflows, the majority of carriers remain focused on foundational capabilities such as data quality, operating model alignment, governance structures, and workforce readiness.

Froseth explained that agentic AI differs from traditional generative AI in that agents are designed not simply to generate content, but to execute tasks and orchestrate workflows by combining large language models (LLMs) with automation and process management tools. These systems can learn from repeated interactions and improve over time. He emphasized, however, that current implementations generally preserve human-in-the-loop controls for higher-risk decisions, particularly in underwriting and claims handling. Examples were shared demonstrating how agents can accelerate operational tasks while maintaining clear human accountability for final determinations.

Froseth and Ilana Golbin Blumenfeld (PwC) stated that change management is an important part of implementation because if employees do not understand the strategy of using AI, they may default to continuing to operate without using the newly available tools.

The presentation also addressed emerging AI operating models within insurance organizations. Froseth described a progression from centralized AI governance structures toward more hybrid approaches as adoption matures. He observed that fully federated models can lead to challenges such as tool proliferation, inconsistent governance, and fragmented risk oversight, while fully centralized models may struggle to scale or meet business demand. Many insurers are now experimenting with hybrid models that balance centralized standards and guardrails with increased flexibility for business units as AI tools become more accessible to non-technical users. Blumenfeld said that centralized operating models sometimes leave units competing for resources, making it challenging to implement projects when teams lack access to expertise. However, generative AI allows business units to do more initial work.

Froseth and Blumenfeld then discussed the evolving risk landscape of agentic AI systems, highlighting new governance challenges, including emergent behavior when multiple agents interact, the potential for cascading failures, automation bias, and increased complexity in managing data access across interconnected systems. She noted that traditional AI governance frameworks—largely developed for static or narrowly scoped models—are insufficient for autonomous or semi-autonomous agent systems.

As a result, insurers are increasingly investing in observability, monitoring, and lifecycle management capabilities to support accountability, resilience, and cost control. Blumenfeld noted that if an individual agent is not designed to access a certain data set, companies need to account for an agent's ability to interact with other agents to access that same data. Thinking about governance frameworks, Blumenfeld said that while the AI principles and strategy a company operates under may not need to change, inventory practices, risk registers, and processes to test, monitor, and track risks may need to be redesigned in light of the advent of agentic AI.

Transitioning to a discussion, Schulz asked how data governance considerations for agentic AI differ from those applicable to more traditional AI systems, and what regulators should be attentive to in practice. Blumenfeld responded that while core data governance principles remain unchanged, agentic systems introduce additional complexity because agents access, combine, and pass information across multiple data sources. She emphasized the importance of clearly defining access permissions, validating data sources, and understanding how agent orchestration could inadvertently circumvent existing controls, requiring more granular oversight than prior AI implementations. Froseth provided an example of working with an AI agent that was given human resources-related tasks, and they worked to ensure access was tiered and granted as appropriate, but thinking about salary information, the agent would have access to.

Commissioner Yaworsky also asked about the importance of responsible AI governance in an environment where insurers face pressure to rapidly deploy AI technologies. Blumenfeld explained that effective governance enables faster, more confident adoption by clarifying expectations for risk tolerance, roles, decision authority, and controls throughout the AI lifecycle. She noted that without such structures, organizations may either move too cautiously or adopt technology without fully understanding risks, potentially leading to consumer harm or operational disruption.

Commissioner Yaworsky said he has heard legislative discussions that tend to center on outcomes, and while that is important, it does not capture the full universe of responsible AI governance, or at least the full scope of a governance framework. Blumenfeld agreed, saying that the insurance life cycle includes many components that are each important, but also cautioned against relying on humans in the loop because of the imperfections that people also have.

Commissioner Hershman raised questions about whether AI governance approaches should be principles-based or prescriptive, with specific deviation thresholds, particularly when performance deviations occur. Blumenfeld stated that rigid, uniform thresholds are difficult to apply across diverse AI use cases and may not adequately reflect contextual risks, compensating controls, or data limitations. She noted that while prescriptive requirements may be appropriate in highly regulated or high-impact contexts, most AI systems require use-case-specific evaluation informed by system design, governance maturity, and the broader control environment, rather than relying on a single quantitative benchmark.

Froseth added that the risk may also vary based on whether the software is third-party software or internally developed. Commissioner Hershman wondered if, at some point, a deviation standard would become necessary. Froseth discussed an example of an AI model's test in which the AI achieved 90% accuracy, but on investigation, the errors were found to be related to human labeling of the data.

Commissioner Samsom asked whether PwC had predictions of AI failures in insurance or more generally, and how governance would respond. Froseth discussed how inappropriate access to production data might pose a risk. Commissioner Samsom then asked whether that might be a relevant consideration in relation to Commissioner Hershman's comments on deviation standards—seeking production data access.

Blumenfeld said she tends to see things optimistically, but remarked that the AI Incident Database provides a public register of AI failures, which may be instructive. In some instances, agent-pushed code has led to system failures, but organizations are thinking critically as they implement AI. Froseth provided an example of Amazon ceasing the use of AI for coding due to system issues AI coding caused.

Commissioner Fowler asked whether the presenters had any recommendations for regulators to consider as they continued their regulatory discourse. Blumenfeld remarked that regulators have a difficult job, as they need to develop and update regulatory frameworks in a rapidly changing environment. Blumenfeld recommended continued engagement with the industry and flexibility, but noted that risks will continue to evolve and may require further adaptation. Froseth talked about the role AI may play in helping transition skills and empowering less experienced staff.

5. Other Matters

Director Dunning then reminded meeting attendees about InsurTech on the Silicon Prairie which will include several Commissioners and a speaker from Open AI in attendance.

Having no further discussion, the Innovation, Cybersecurity, and Technology (H) Committee adjourned.

SharePoint/NAIC Support Staff Hub/Committees/H CMTE/2026_Spring/H-Minutes/Minutes-H-Cmte032526-Final.docx

2. Consider Adoption of the Reports of its Working Groups and Subgroup

Attachment Two

Commissioner Michael Yaworsky (FL)

Insurtech

ON THE SILICON PRAIRIE

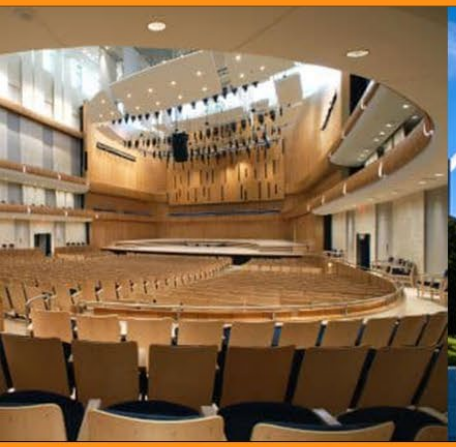

3. Hear a Presentation from Microsoft on Artificial Intelligence (AI) Trends

Attachment Three

Jim DeMarco (Microsoft)

Microsoft

Microsoft Frontier Company

What's the real frontier of AI in insurance?

Jim DeMarco
Insurance Advisory Lead
August 2026

Seismic Changes in Insurance

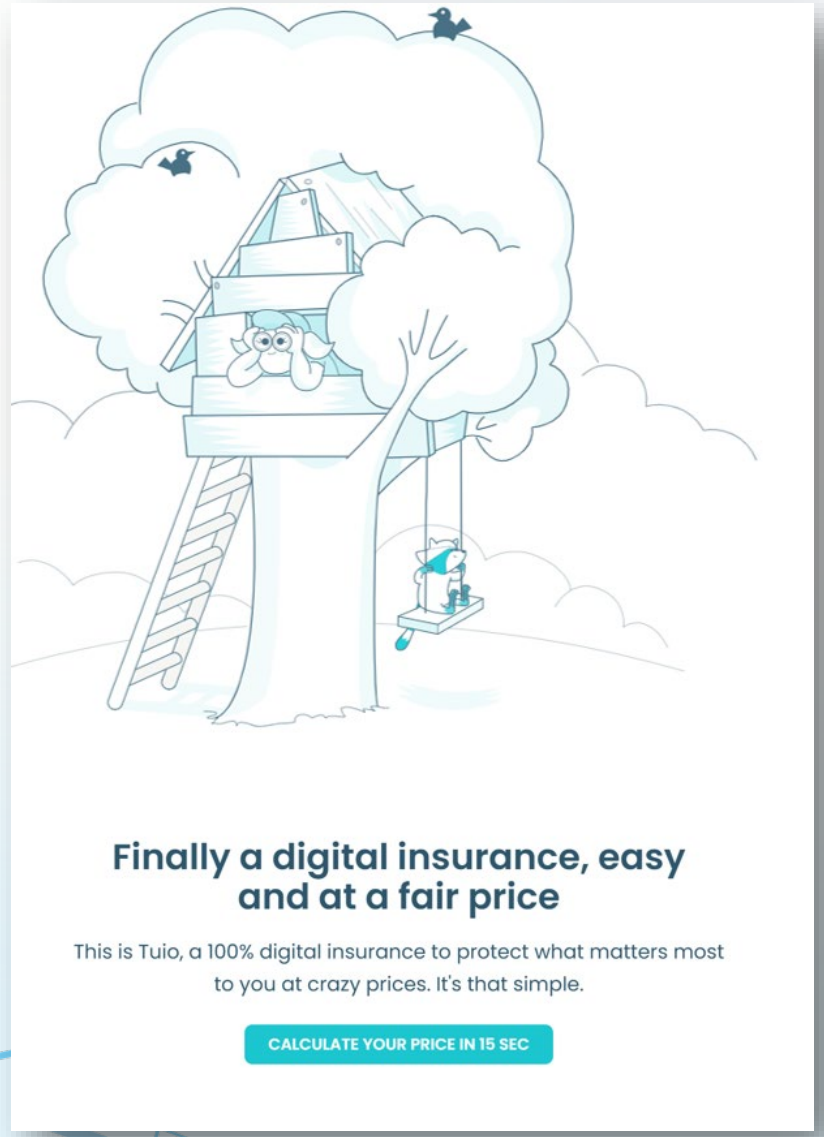

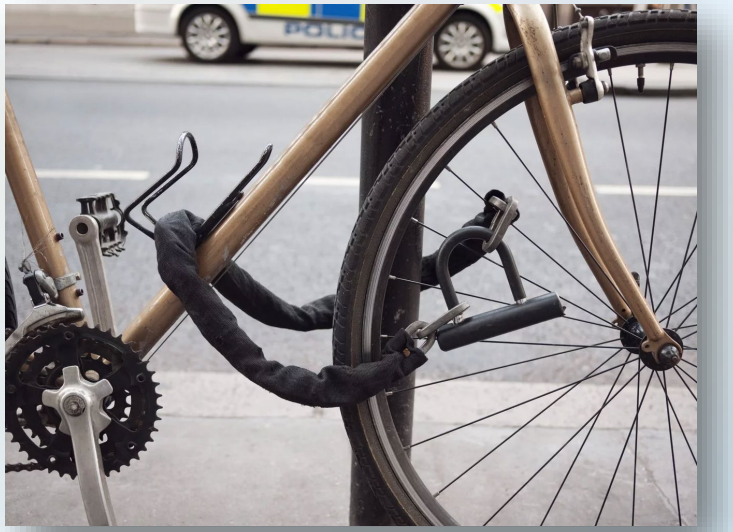

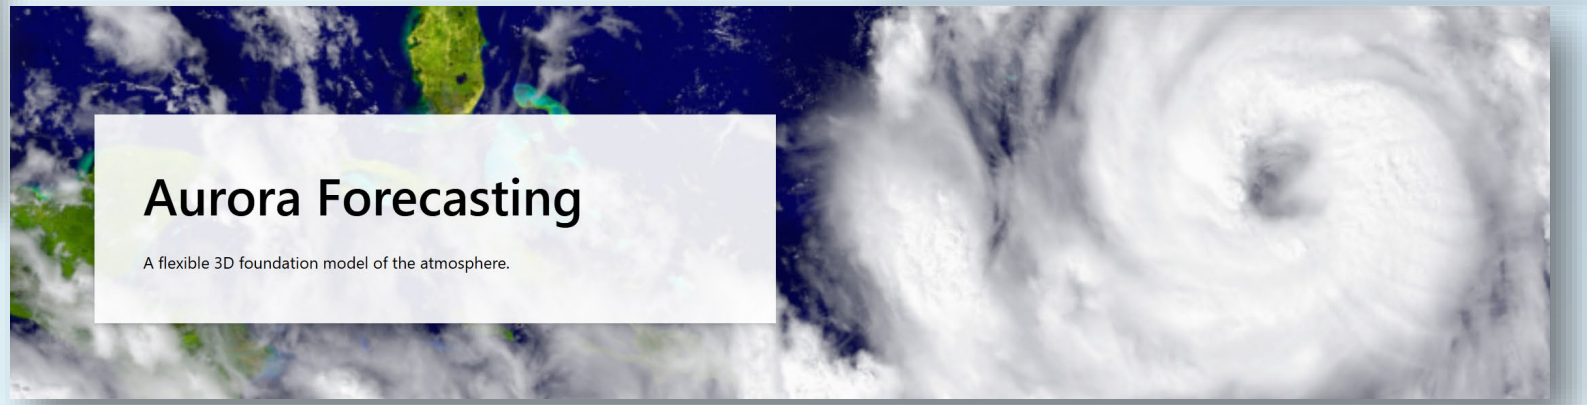

AI in Insurance – What We're Seeing

Agentic AI will transform the industry, but the industry needs help to change.

LLMs are not a differentiator; how you use, reuse, secure, and govern them is.

Agentic AI changes people's roles, it doesn't replace the need for them.

Scaled AI requires systemic intelligence and trust.

AI Adoption Patterns in Insurance

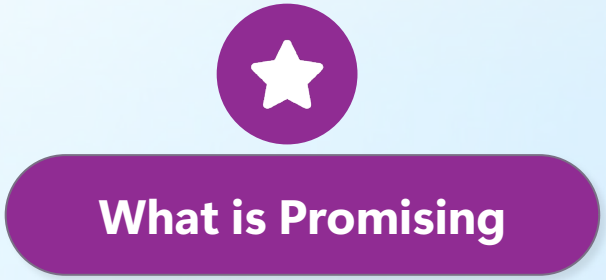

From the new "table stakes" to the overhyped and the unnecessary

AI Adoption Patterns in Insurance

What is Proven

- Internal and external compliant search
- Contextualized document processing
- Insurance professional workbenches
- Underwriting assistants
- Broker/agent assistants
- Call center interaction enhancement and policy enforcement

AI Adoption Patterns in Insurance

What Works

- Guided customer interactions / AI-driven customer guidance
- Straight through underwriting
- Straight through claims processing (HITL for adverse decisions)
- Regulatory filings explainers, formatting automation
- Risk placement automation
- Adaptive compliance guidelines
- Proactive client outreach

AI Adoption Patterns in Insurance

What is Promising

- Emulating keyboarding
- Actuarial modeling using new data sources
- Autonomous underwriting and placement
- Goal-based scenario planning
- Automated risk mitigation

AI Adoption Patterns in Insurance

What is Not Ready

- Human level decision-making by AI
- Fully autonomous claims

To Scale AI Successfully, Governance Must Evolve to Continuous, System-Level Oversight

01

Ecosystem Governance

Financial institutions must govern interconnected ecosystems of models, agents, vendors, data flows, and operational dependencies — not just individual AI models.

02

Growing Systemic Risk

While limited evidence of AI-driven systemic risk exists today, risk is expected to grow as AI moves into critical workloads with greater speed, scale, and autonomy.

03

Oversight at Scale

The central governance challenge is whether oversight, resilience, and control frameworks can keep pace with multi-model and agentic AI environments.

04

Beyond Human-in-the-Loop

Traditional human-in-the-loop oversight will not scale alone; firms need risk-based models supported by documentation, monitoring, and AI-augmented supervision.

Trusted AI as an Enterprise Control Architecture for Safe Adoption at Scale

01

End-to-End Visibility

Trust in AI depends less on one-time validation and more on visibility into reasoning steps, decision chains, outputs, intermediate actions, and operational dependencies.

02

Agents as Enterprise Actors

Agentic AI introduces a new control challenge: agents may need identity, access controls, inventories, monitoring, and behavioral observability — much like enterprise actors.

03

Strategic Opportunity

Responsible AI, security, compliance, and resilience capabilities can be positioned as core enablers of safe AI adoption in regulated industries.

04

Governance Tooling

Capabilities such as Agent 365, Entra, Purview, Foundry, Defender, and broader governance tooling support a customer-ready control narrative for AI risk management.

Data, People and Trust are Converging

More data → more trust → more intelligence.

Intelligence + Trust is a business-model prerequisite - not a technology choice.

POWERED BY ECOSYSTEM INTELLIGENCE

SIGNALS

- Employee feedback
- Operational data
- Customer signal
- Product telemetry

EXTERNAL DATA

- Macro & enriched data (IoT)
- Regulatory data
- Fraud & cyber intelligence
- External data sets

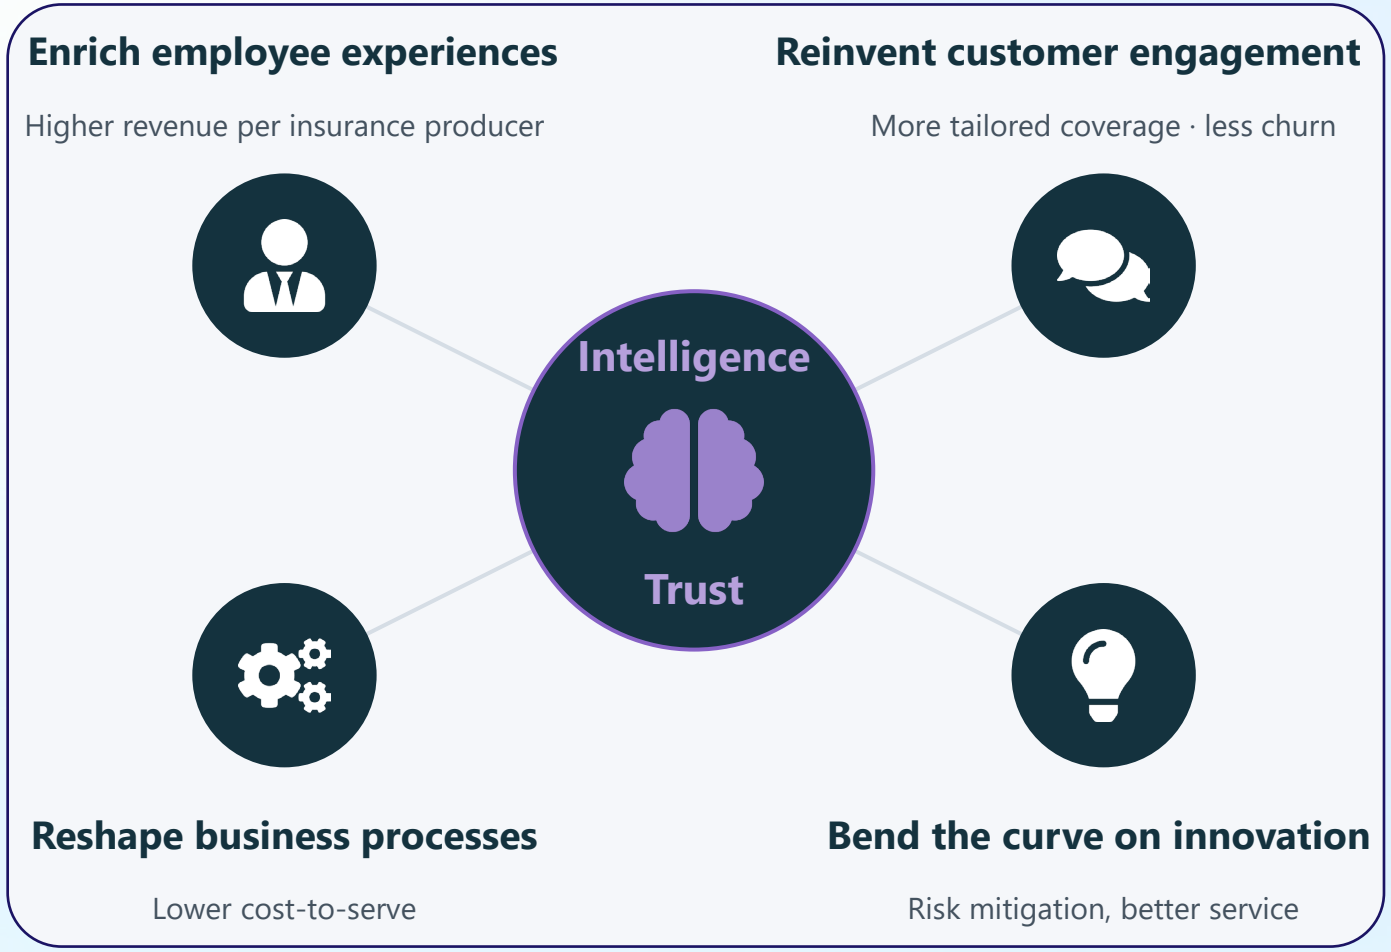

Microsoft

Thank you!

Alignment of Microsoft's Principles to Global Regulatory Intent

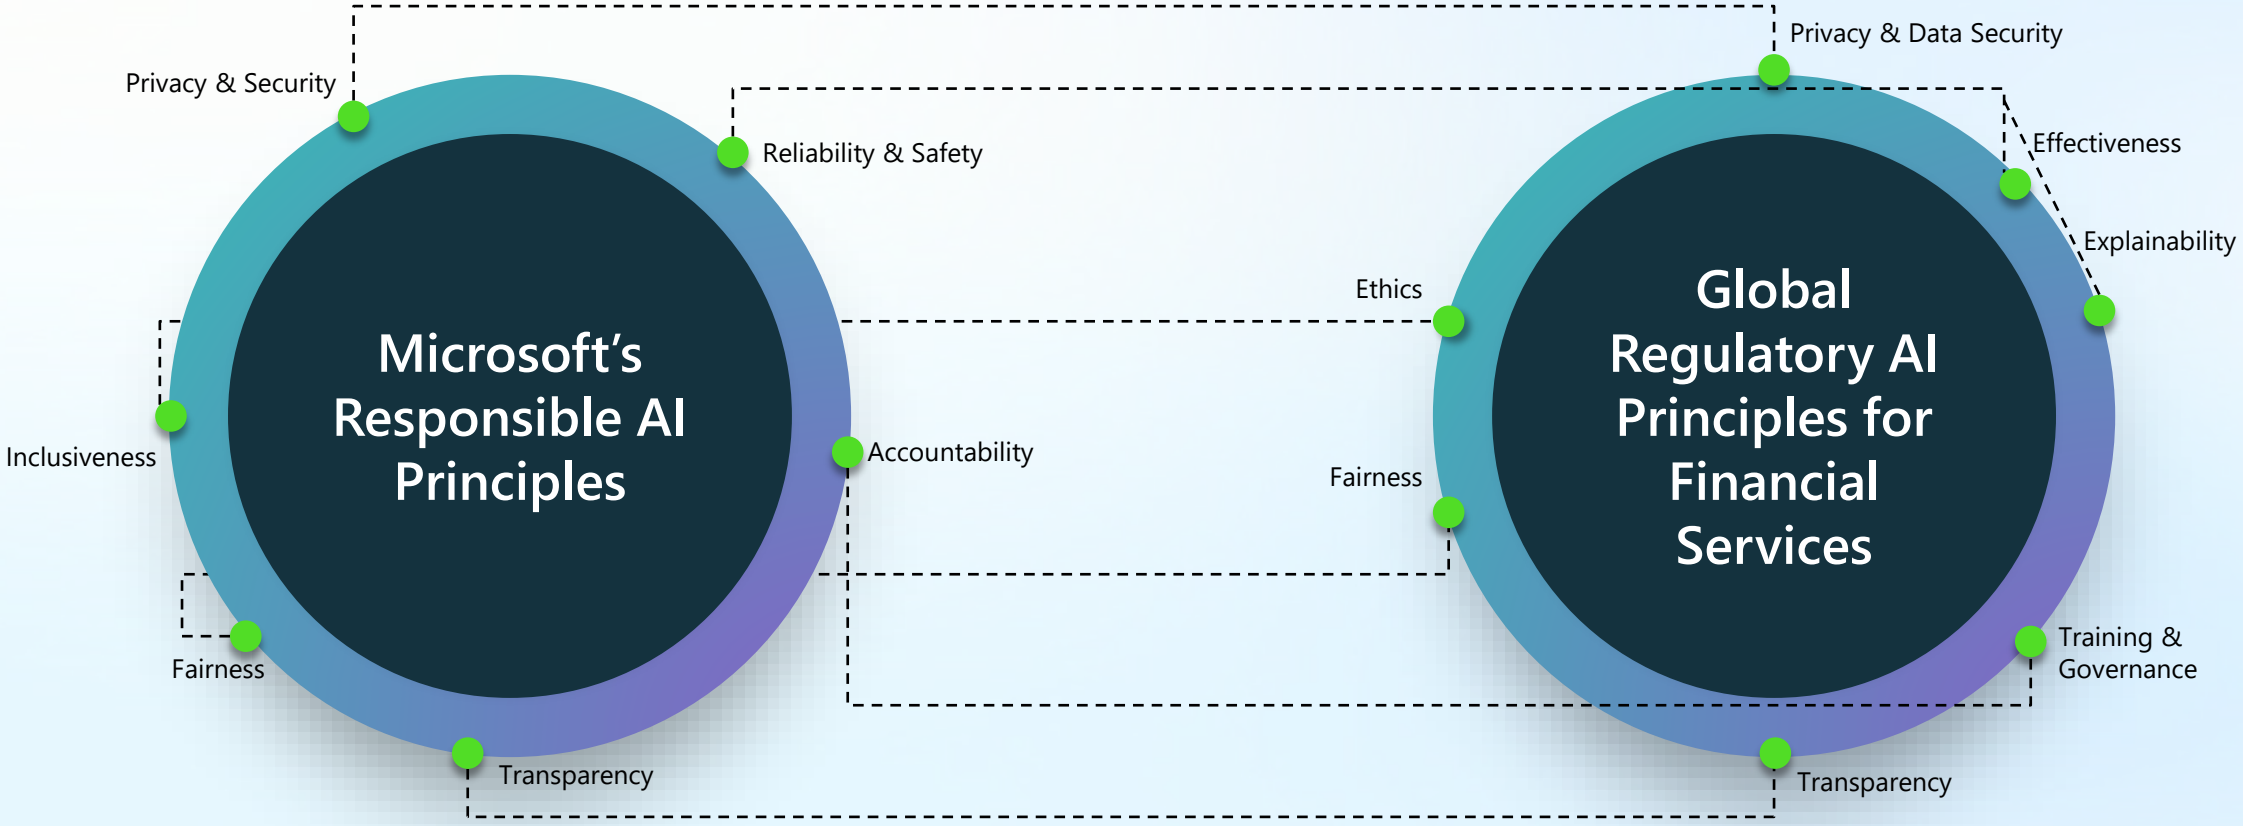

4. Receive an Update on the Cybersecurity Event Notification Portal

Commissioner Michael Yaworsky (FL)