Draft date: 11/10/26

2026 Fall National Meeting
Dallas, Texas

INNOVATION, CYBERSECURITY, AND TECHNOLOGY (H) COMMITTEE
Tuesday, November 17, 2026
12:15 – 1:30 p.m.
Convention Center—Texas AB—Level 3

ROLL CALL

Michael Yaworsky, Chair	Florida	Angela L. Nelson	Missouri
Karima M. Woods, Vice Chair	District of Columbia	James E. Brown	Montana
Mark Fowler	Alabama	Eric Dunning	Nebraska
Heather Carpenter	Alaska	Jon Godfread	North Dakota
Jimmy Harris	Arkansas	Michael Humphreys	Pennsylvania
Michael Conway	Colorado	Carter Lawrence	Tennessee
Doug Ommen	Iowa	Kaj Samsom	Vermont
Marie Grant	Maryland		

NAIC Committee Support: Miguel Romero/Scott Morris/Scott Sobel

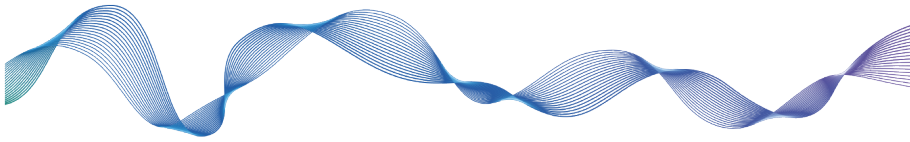
AGENDA

1. Consider Adoption of its April 30 and Summer National Meeting Minutes—*Commissioner Karima Woods (DC)* — Attachment One

2. Consider Adoption of the Reports of its Working Groups—*Commissioner Karima Woods (DC)* — Attachment Two
   A. Big Data and Artificial Intelligence (H) Working Group
   B. Cybersecurity (H) Working Group
   C. Privacy Protections (H) Working Group
   D. SupTech/GovTech (H) Subgroup
   E. Third-Party Data and Model (H) Working Group

3. Consider Adoption of the AI Risk Evaluation Supplement —*Commissioner Michael Yaworsky (FL)* — Attachment Three

4. Consider Adoption of its 2027 Proposed Charges —*Commissioner Michael Yaworsky (FL)* — Attachment Four



5. Hear a Presentation on How the Audit Profession is Adapting for Artificial Intelligence (AI)—*Michelle Donahue (Deloitte)* — Attachment Five

6. Discuss Any Other Matters Brought Before the Committee —*Commissioner Michael Yaworsky (FL)*

7. Adjournment

Draft Pending Adoption

Draft: 8/27/26

Innovation, Cybersecurity, and Technology (H) Committee
Columbus, Ohio
August 14, 2026

The Innovation, Cybersecurity, and Technology (H) Committee met in Columbus, OH, August 14, 2026. The following Committee members participated: Michael Yaworsky, Chair (FL); Heather Carpenter (AK); Mark Fowler (AL); Jimmy Harris (AR); Michael Conway and Jason Lapham (CO); Doug Ommen and Daniel Mathis (IA); Marie Grant (MD); James E. Brown represented by Erin Snyder (MT); Jon Godfread (ND); Eric Dunning (NE); Michael Humphreys represented by Jodi Frantz and Michael McKenney (PA); Carter Lawrence represented by Scott McAnally (TN); and Kaj Samsom and Mary Block (VT). Also participating were: Joshua Hershman and George Bradner (CT); Judith L. French (OH); Glen Mulready (OK); Elizabeth Kelleher Dwyer and Matthew Gendron (RI); Amanda Crawford and Marianne Baker (TX); Jon Pike (UT); and Scott A. White, Eric Lowe, and Michael Peterson (VA).

1. Adopted the Reports of its Working Groups and Subgroup

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

Block reported that the Big Data and Artificial Intelligence (H) Working Group met Aug. 13. During that meeting, the Working Group adopted its July 22 minutes. The Working Group then received an update on the Artificial Intelligence Risk Evaluation Supplement (AIRES) pilot process. The pilot, which began in March, involves 12 participating states and will continue through September. Through the pilot, states are learning about the use of artificial intelligence (AI) while providing feedback that will help clarify many aspects of the supplement. The next version of the supplement will be shared at the end of August, followed by two rounds of public exposure. Public updates will be provided throughout the pilot process, with the goal of adopting the supplement at the Fall National Meeting.

The Working Group closed its meeting by hearing a presentation from Edin Imsirovic (AM Best) on AI governance in insurance. The presentation defined predictive AI, generative AI, and agentic AI; described innovation and enterprise risk management (ERM) perspectives for evaluating governance quality and how governance can affect value and risk; and discussed the inputs, outputs, and control challenges across AI capability tiers.

B. Cybersecurity (H) Working Group

Peterson reported that the Cybersecurity (H) Working Group met June 2 in regulator-to-regulator session, pursuant to paragraph xx (insert appropriate text) of the NAIC Policy Statement on Open Meetings, to demonstrate the proposed cybersecurity event notification portal. Kim Myers (NAIC) demonstrated a nonfunctional prototype to illustrate the tool's intended capabilities across both an industry licensee interface and a regulator interface. The tool supports structured question-and-answer workflows at the event level, with resolution tracking, reminders, and exportable records comparable to existing Uniform Certificate of Authority Application (UCAA) functionality. The portal proposal has been reviewed and approved at the Working Group, Committee, and Enterprise Project Management Office (EPMO) levels, and the next step is for the Executive (EX) Committee to review and approve it.

The Working Group then met June 10 and heard a presentation on the NetDiligence 2025 *Cyber Claims Study*. The study draws on a five-year data set of about 10,400 claims focused primarily on small- and mid-sized enterprises (SMEs). According to the study, the average SME incident cost was approximately $203,000 in 2024, compared with approximately $5 million for large entities. Ransomware was cited as a leading peril and the top driver of

© 2026 National Association of Insurance Commissioners

demand, with SME demands averaging $3.1 million and then negotiated down to about $1 million. A notable shift from prior years was that nearly all 2024 claims were tied to criminal activity, with business email compromise identified as the second-leading cause of loss, typically involving automated clearing house (ACH) or wire fraud through executive impersonation. In discussion, the Working Group explored the data elements needed to move from tracking claim volume toward understanding incident impacts and the frequent inadequacy of purchased coverage limits relative to actual incident costs. The Working Group is looking to coordinate multiple interim meetings in the coming months to hear about and discuss the emerging risks around quantum computing and post-quantum cryptography, as well as the impact of frontier AI models on the cybersecurity landscape.

C. Privacy Protections (H) Working Group

Director Dwyer reported that the Privacy Protections Drafting Group met June 3 and April 30 to consider changes to Article 7 and Article 8 of the *Privacy of Consumer Financial and Health Information Regulation* (#672).

On July 24 the Working Group exposed the full draft of Model #672, incorporating all of the work completed on the individual sections, for a 60-day public comment period ending Sept. 22.

The Working Group also Aug. 13 to hear verbal comments on the full draft of the revised Model #672. Once the written comments are received, the Working Group will compile and review them. After the chairs review the comments, the Working Group will communicate its next steps, which will depend on the comments received.

D. SupTech/GovTech (H) Subgroup

Bradner reported that the SupTech/GovTech (H) Subgroup focuses on regulatory tools and education and hears from companies on their proprietary technology. Its meetings are conducted in regulator-to-regulator session following the NAIC Policy Statement on Open Meetings.

The Subgroup continues to use the results of a regulator survey conducted last year to identify topics that drive efficiency, reduce administrative burden, and strengthen regulatory effectiveness. Since the Spring National Meeting, the Subgroup met twice. On May 27, it convened state insurance regulators to exchange ideas on how insurance departments are approaching AI adoption and the use of AI tools in insurance regulation, including practical lessons learned, challenges, and open questions. On July 28, the Subgroup viewed a demonstration of Deloitte's resiliency navigator system, a consumer-facing platform that helps state departments of insurance (DOIs) close the consumer education gap before and after a loss. The Subgroup continues to evaluate additional presentations and educational opportunities for 2026 on technical topics that would benefit regulators' collective understanding and improve efficiency and oversight capabilities.

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

Lapham reported that the Third-Party Data and Models (H) Working Group met Aug. 12. During that meeting, the Working Group adopted its July 16 and Spring National Meeting minutes and heard comments on the exposed draft third-party regulatory framework for property/casualty (P/C) pricing and underwriting. He said that Working Group received 11 comment letters from interested parties on the exposed draft framework. The general themes of the comments included: 1) concerns with the attestation asking vendors to certify compliance they cannot control; 2) inadequate and uncertain confidentiality protections for vendor filings across states; 3) whether registration would be voluntary given the consequences of use restrictions; 4) the need to define material change objectively; and 5) duplication of existing state filing and oversight processes. The drafting group will meet soon to weigh revisions to the draft framework based on this recent feedback for consideration by the full Working Group, with the hope that the framework will be ready for potential adoption at the Fall National Meeting.

Additionally, it was reported that the Committee met April 30 (Attachment One). During this meeting, it took the following action: 1) adopted its Spring National Meeting minutes (*see NAIC Proceedings – Spring 2026, Innovation, Cybersecurity, and Technology (H) Committee*); and 2) adopted the cybersecurity event notification portal project proposal.

Commissioner Carpenter made a motion, seconded by Commissioner Conway, to adopt the reports of the: Big Data and Artificial Intelligence (H) Working Group (Attachment Two), including its July 22 and June 1 minutes; Cybersecurity (H) Working Group (Attachment Three), including its June 10 minutes; Privacy Protections (H) Working Group (Attachment Four); Third-Party Data and Models (H) Working Group (Attachment Five), including its Feb. 26 minutes; SupTech/GovTech (H) Working Group; and Data Call Study Group. The motion passed unanimously.

2. Heard a Presentation from Microsoft on AI Trends in the Insurance Industry

The Committee heard a presentation from Microsoft on AI trends in the insurance industry. Jim DeMarco (Microsoft), an insurance advisor lead, outlined four broad trends reshaping the industry.

DeMarco first pointed to automation, noting that the first insurance claim processed solely by AI occurred in 2023—a stolen-bicycle claim in the United Kingdom (UK) that took two seconds to process. He next described a shift in distribution, observing that the first policy offered, bound, and sold entirely through ChatGPT occurred earlier this year in Spain, after which the small insurer that offered the policy, Twee O, was acquired by a larger competitor. Third, he addressed demographic change, citing reporting on a study suggesting a maximum human age of 194 years and noting that people are living longer but not necessarily healthier, with implications for wellness, elder care, financial risk transfer, and the broader wealth transfer. Fourth, he discussed weather and weather modeling, noting that severe convective storms are becoming more frequent and severe, that Tornado Alley has shifted east. He said that AI is improving both longer-range forecasting and imminent-risk identification, creating opportunities for risk mitigation and prevention services that lower the cost of loss. DeMarco said these trends are changing the fundamental assumptions underlying how the industry views risk, risk transfer, and claims payment.

Turning to the use of AI in the insurance industry, DeMarco focused on agentic AI—the granting of agency to predictive models to make decisions—which he said could transform underwriting, claims, and distribution. However, he cautioned that AI is a tool rather than an inherently intelligent system and must be adapted to the appropriate use case. He said large language models (LLMs) from various providers are not themselves a differentiator. Rather, the differentiator is how they are used, and a tool that improves remains ineffective in the hands of someone using it poorly. He said AI is changing jobs rather than eliminating the need for human workers, using the analogy that stronger hurricane-modeling tools narrow predicted storm paths but do not remove the need for catastrophe modelers. On scaling, DeMarco cited a Boston Consulting Group report ranking insurance second only to the technology industry in AI use but near the bottom in scaling that use, observing that the industry tends to experiment and leave capabilities in isolated pockets. He said scaling AI across an organization requires systemic governance that brings intelligence to the point of decision while ensuring its use is trustworthy.

DeMarco organized current use cases into four categories. He described as proven, or table stakes: 1) the use of compliant grounded search with natural language processing and generative AI to reduce response times from days to seconds; 2) contextualized document processing; 3) guidance to professionals in claims and underwriting, including underwriter-assistance tools that are more mature outside the U.S.; and 4) assistance for brokers, agents, and call-center interactions.

In terms of what works, DeMarco said the greatest current energy is in guiding customer interactions and in straight-through processing for underwriting, while similar automation for claims is advancing but not fully realized, particularly for adverse decisions. For those decisions, DeMarco said Microsoft advocates keeping humans in the loop. He cited the analogy that although autopilot has existed for flight and landing since 1947 and for takeoff since 2024, passengers still want a pilot on board. Predictive AI is very good at determining what could happen next, but it is not suited to deciding what should happen.

Among promising developments, DeMarco pointed to tools that emulate a person at a keyboard to work with legacy mainframe systems, tools that identify new data to test actuarial assumptions, and tools that automate risk mitigation, such as increasingly precise alerts warning policyholders of approaching hailstorms.

DeMarco said AI is not ready to provide judgment, because predictive AI—including generative and agentic AI—lacks abductive reasoning, so human-level decision-making remains necessary.

On governance, DeMarco said that scaling AI immediately raises governance questions and that financial institutions, including insurers, must govern their entire interconnected ecosystems and data—both their own and data received from others—rather than individual models alone. He noted the systemic risk that governance itself can pose, citing the example of a model that could no longer be exported outside the U.S., and said oversight must occur at the organizational level. He described a framework for trust that is both architectural and systemic, in which AI agents have identities, are restricted in what they can access as any employee would be, and can explain what they have done and how they used data, supported by a responsible AI framework. He concluded that AI is increasingly providing intelligence at the point of decision and that failing to manage trust would provoke backlash from the public and customers alike.

Following the presentation, Commissioner Yaworsky described a scenario in which a health insurer compared its usual human-involved claims process with a nearly fully automated AI process that produced more accurate and better outcomes. However, he said most participants in his discussions still preferred substantial human involvement even when outcomes were poorer, reflecting a desire for humanity and trust in the process. DeMarco responded that an auto insurer that shifted first notice of loss to AI achieved strong call-deflection rates and responsiveness, but 35% of customers immediately requested a representative because people under stress want to speak with people. He noted that studies indicate that human decision-makers tend to reach better outcomes earlier in the day than later. He said the key question is the right outcome for the claimant—such as getting a car back on the road or access to a medication—rather than merely processing a claim, adding that simply digitizing an existing human process is unlikely to improve outcomes.

Miguel Romero (NAIC) said that state insurance regulators had heard in a prior session about the growing use of AI to generate consumer complaints, which can make it harder to identify the underlying issue and verify the facts. He asked whether solutions could help consumers advocate for themselves without creating additional workload for companies or regulators. DeMarco said that every tool can be used as a weapon and that some bad actors may be using AI to submit high volumes of complaints. He drew a parallel to the evolution of social media and to cybersecurity, where AI is used by criminals and state-sponsored actors, and the response must likewise incorporate AI to predict, identify, and respond. He said insurance departments will likely need tools to distinguish genuine complaints from fabricated ones and then respond appropriately—human-to-human contact when needed, human-to-technology contact for information seeking, and technology-to-technology filtering for spam. DeMarco said he anticipates the increasing use of personalized agents that will, over time, communicate with regulators and insurers on consumers' behalf.

© 2026 National Association of Insurance Commissioners

3. Received an Update on the Cybersecurity Event Notification Portal

Commissioner Yaworsky provided an update on the cybersecurity event notification portal. He noted that the NAIC had experienced a cybersecurity event several weeks earlier and said the purpose of the discussion was not to review that event but to acknowledge and address concerns raised afterward about the NAIC creating a portal and the data that would be housed within it.

Commissioner Yaworsky recalled that the proposal to adopt the portal project had been adopted at the Spring National Meeting. Before the project is considered by the NAIC's Executive (EX) Committee, he asked NAIC committee support to work with Commissioner Crawford to perform a security review of the project, noting that her background in technology would help ensure the membership could move forward responsibly. He said this was a novel situation without a predefined scope for the review, but he expressed confidence in Commissioner Crawford's ability to guide the effort. He suggested that one component could be a clear, easy-to-read listing of the security commitments the NAIC provides to its membership and regulated entities so everyone could understand the safeguards in place for data housed in the portal. He said much of this had already been discussed during the portal's development and expressed hope that the review would be brief. Commissioner Yaworsky invited questions or comments, and none were raised.

4. Discussed Other Matters

Director Dunning announced that Commissioner Crawford, Commissioner Godfread, Commissioner Yaworsky, and Director Dwyer are scheduled to speak Oct. 20 at Insurtech on the Silicon Prairie, which will be held Oct. 20–21 in Omaha, NE. Having no further business, the Innovation, Cybersecurity, and Technology (H) Committee adjourned.

SharePoint/NAIC Support Staff Hub/Committees/H CMTE/2026_Summer/H-Minutes/Minutes-H-Cmte081426-Final.docx

© 2026 National Association of Insurance Commissioners

Draft Pending Adoption

Attachment One
Innovation, Cybersecurity, and Technology (H) Committee
8/14/26

Draft: 8/27/26

Innovation, Cybersecurity, and Technology (H) Committee
Virtual Meeting
April 30, 2026

The Innovation, Cybersecurity, and Technology (H) Committee met virtually on April 30, 2026. The following Committee members participated: Michael Yaworsky, Chair (FL); Karima M. Woods, Vice-Chair (DC); Heather Carpenter and Molly Nollette (AK); Mark Fowler represented by Richard Fiore (AL); Jimmy Harris represented by Chris Erwin (AR); Michael Conway represented by Jason Lapham (CO); Doug Ommen (IA); Marie Grant (MD); Angela Nelson (MO); Jon Godfread represented by Colton Schulz (ND); Eric Dunning (NE); Carter Lawrence (TN); and Mary Block and Jared Holshouser (VT). Also participating were: Scott A. White and Michael Peterson (VA)

1. Adopted its 2026 Spring National Meeting Minutes

Schulz made a motion, seconded by Commissioner Lawrence, to adopt the Committee's March 25, 2026, minutes (*see NAIC Proceedings – Spring 2026, Innovation, Cybersecurity, and Technology (H) Committee*). The motion passed unanimously.

2. Adopted the Cybersecurity Event Notification Portal Project Proposal

The Committee next considered the Cybersecurity Event Notification Portal Project Proposal. Commissioner Yaworsky noted the portal would allow states to receive cybersecurity event notices through an NAIC-developed application, simplify compliance burdens, enhance the security of notification information, and leverage the NAIC's capacity to develop in-house applications tailored to regulator input. The Chair stated the proposal had been discussed extensively at both the Committee and working group levels and that the Cybersecurity (H) Working Group adopted the project proposal at its March 13, 2026, meeting.

Peterson provided an overview of the project. Peterson described the project's origin following the adoption of the *Insurance Data Security Model Law* (#668) and the Cybersecurity Event Response Plan. He explained that at the 2024 Fall National Meeting, the NAIC adopted a motion directing NAIC committee support to explore the creation of the portal. In 2025, NAIC committee support developed detailed project documentation, which underwent multiple rounds of public comment from both regulators and industry. Following the incorporation of those comments, the project documentation reached a stable consensus and was adopted by the Working Group in March 2026.

Peterson stated the portal is intended to address fragmentation in cybersecurity event notification requirements across the 28 jurisdictions that have passed versions of Section 6(B) of Model #668. He noted the portal is designed to centralize and standardize the notification process by providing a single standardized form that captures all jurisdictional variations. He stated this approach would reduce current and future compliance costs for licensees. Peterson further noted that for states that have not yet adopted the model law, the availability of the portal could serve as an incentive, because adopting the law in conjunction with the portal would effectively reduce rather than increase regulatory burden on licensees. Peterson described the confidentiality and security features of the portal, noting that the NAIC agreed to provide a Service Organization Control (SOC) 3 report to the industry as an assurance mechanism on the NAIC's security.

© 2026 National Association of Insurance Commissioners

Draft Pending Adoption

Yaworsky provided contextual remarks, referencing prior large-scale cybersecurity breaches affecting the health industry and the challenges regulators faced coordinating responses across jurisdictions without a unified notification system. The Chair stated that a unified portal would significantly improve regulators' ability to assess severity and take timely action. He commended Peterson and the working group for navigating a technically complex and lengthy process.

Commissioner Grant made a motion, seconded by Lawrence, to adopt the Cybersecurity Event Notification Portal Project Proposal (Attachment One-A). The motion passed unanimously. The Chair stated the proposal would be forwarded to the NAIC Executive Committee for consideration.

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

SharePoint/NAIC Support Staff Hub/Committees/H CMTE/2026_Summer/2026 0430Interim-Meeting/Minutes-H-Cmte043026-Final.docx

© 2026 National Association of Insurance Commissioners

2026 Fall National Meeting
Dallas, Texas

BIG DATA AND ARTIFICIAL INTELLIGENCE (H) WORKING GROUP
Monday, November 16, 2026
12:00 p.m. – 1:00 p.m.
Convention Center—Grapevine CD—Level 3

Meeting Summary Report

The Big Data and Artificial Intelligence (H) Working Group met Nov. 16, 2026. During this meeting, the Working Group:

1. Adopted its Nov. 2 minutes. During this meeting, the Working Group took the following action:
   A. Adopted its Oct. 8 minutes. During this meeting, the Working Group took the following action:
      i. Adopted its Aug. 31 meeting minutes. During this meeting, the Working Group took the following action:
         a. Adopted its Summer National Meeting minutes (*see NAIC Proceedings – Summer 2026, Innovation, Cybersecurity, and Technology (H) Committee, Attachment Two*).
         b. Received an update on the proposed changes to the Artificial Intelligence (AI) Risk Evaluation Supplement.
      ii. Received an update on the supplement pilot process.
      iii. Heard comments from interested parties summarizing their input on version 5.0 of the supplement.
   B. Received an update on the supplement pilot process.
   C. Heard comments from interested parties summarizing their input on version 6.0 of the supplement.

2. Received an update on the supplement pilot process. The pilot was conducted in 12 participating states from March through September. The Working Group held several public meetings, including at the Summer National Meeting, released version 5.0 in September and version 6.0 in October, and received extensive written and verbal comments. Version 7.0 was released Nov. 9, and the Working Group heard additional verbal comments during its Nov. 16 meeting.

3. Adopted the supplement after reaching a consensus to agree on the two technical edits raised.

Draft: 11/9/26

Cybersecurity (H) Working Group
Virtual Meeting
November 5, 2026

The Cybersecurity (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met Nov. 5, 2026. The following Working Group members participated: Michael Peterson, Chair (VA); Colton Schulz, Vice Chair (ND); Alex Romero and Julia Jette (AK); Mark Fowler (AL); Chris Erwin and Hank House (AR); Lori Dreaver Munn (AZ); Damon Diederich (CA); Wanchin Chou and Christian Citarella (CT); Tim Li (DE); Matt Kilgallen, Elizabeth Nunes, and Paula Shamburger (GA); Daniel Mathis (IA); Shane Mead (KS); Nina Hunter (LA); Mary Kwei (MD); Danielle Torres, Alex Hertrich, and Johnny Martin (MI); Gregory Maus and Bubba Aguirre (MN); Kim Dobbs (MO); Troy Smith (MT); Tracy Biehn (NC); Martin Swanson (NE); Joshua Hilliard (NH); Michael Lord (NM); Gille Ann Rabbin (NY); John Pollock and Matt Walsh (OH); David Buono (PA); Curt Kwak (WA); Jamie Adams, Barbara Belling, Adam Donovan, Jeff Grothman, John Litweiler, and Matt Bradley (WI); and Lela D. Ladd (WY). Also participating were: Brian Fjeldheim, John Haworth, and Joshua Blakey (OR); Matt Gendron (RI); Tony Dorschner (SD); and Tracy Klausmeier (UT).

1. Received an Overview of the 2025 Cyber Supplement Data and Key Findings from the Annual Cybersecurity Insurance Market Report

Peterson welcomed members and interested regulators and introduced Koty Henry (NAIC), Cybersecurity Policy Advisor. Peterson characterized the presentation as a market update and regulator discussion rather than a request for an immediate decision on new reporting requirements. He said the discussion would address the market, threat landscape, use of the report's findings, strengths and limitations of current data, and potential next steps. Henry presented the 2025 Cyber Supplement data and findings from the 2026 Report on the Cybersecurity Insurance Market.

Henry reported $9.46 billion in total U.S. direct written premium, including alien surplus lines, an increase of 3.5% from the prior year. Domestic insurers wrote $7.46 billion. Policies in force increased by 33% to 5.84 million, while reported claims increased by 20% to 56,323. Henry said premium growth was modest compared with the more rapid growth in exposures and claims.

Henry reported that the industry loss and defense-and-cost-containment-expense ratio increased from 48.7% to 53%, the second consecutive annual increase and the highest level since the COVID-era ransomware surge. At the same time, U.S. cyber rates continued to decline. He said the combination of rising losses and falling prices differs from the 2020–2021 ransomware cycle and could test rate adequacy if the pattern continues, although the market outlook remains stable.

Henry described distinct admitted and surplus lines segments. Surplus lines insurers held nearly two-thirds of premium, primarily through standalone primary and excess cyber policies, while the admitted market consisted largely of cyber endorsements to other commercial policies. The surplus lines loss ratio was approximately 55.9%, compared with 50.2% for admitted insurers. Third-party claims increased by almost 40% in surplus lines, compared to approximately 12% in the admitted market.

Henry said the market remains concentrated, with the top 20 groups writing approximately 77% of premium and the top 10 writing approximately half of the market. He identified Chubb, Beazley, and Travelers as leaders in the primary segment. He said recent growth partly reflected movement of business from offshore entities to U.S. entities and that recent acquisitions would further affect concentration.

© 2026 National Association of Insurance Commissioners

Citing a NetDiligence study of more than 10,000 claims, Henry said large companies accounted for approximately 3% of claims but 56% of total incident costs. For small businesses, claims involving business interruption cost more than five times as much as claims without business interruption. Insurance paid approximately two-thirds of incident costs for small businesses but approximately one-fifth for large companies. Henry noted that costs absorbed through retentions, sublimits, exclusions, or lack of insurance are not captured by the Cyber Supplement.

Henry reviewed the changing threat environment. He reported that during the first half of 2026, 88% of observed exploitation occurred within 48 hours after public exploit code was released, zero-day exploitation increased by 42%, and artificial intelligence (AI)-enabled adversary activity increased by 89%. He said the Cyber Supplement provides an annual, retrospective view, while cyber risk can change within hours.

Henry identified identity and access management weaknesses, social engineering, and fraud as significant sources of claims. He reported that voice-phishing intrusions doubled, device-code phishing increased fifteenfold, and business email compromise claim frequency increased by approximately 30% in late 2025, primarily among smaller businesses. He said funds-transfer fraud losses were forecast to increase by 20%–50%. A typical $250,000 sublimit may cap loss severity but does not reduce claim frequency, and fraud losses can fall under cyber or crime coverage, complicating analysis.

Henry identified correlated losses arising from shared technology as a defining challenge. He said one threat group compromised more than 300 software dependencies in a single day, cloud-focused criminal activity increased by approximately 171%, and supply-chain business interruption increased from 6% to 15% of large-claim value in one year. A single event involving a widely used provider can affect many policyholders through both affirmative and silent cyber exposures, while available data for modeling such events remains limited.

Henry also discussed growing liability exposure. He said privacy claims doubled among one insurer's policyholders during the first half of 2026. Claims that do not involve a data breach, including claims related to tracking pixels, privacy, biometrics, and AI profiling, are also increasing and may overlap cyber, directors and officers, and errors and omissions coverage. Henry said increasing third-party claims and class actions may raise social inflation risk and noted that third-party claims generally take longer to develop.

Henry described a continuing protection gap. He said approximately 1% of global economic cyber losses are insured and that surveys suggest awareness and distribution, rather than price alone, contribute to the gap. Approximately one-quarter of surveyed executives did not know cyber insurance existed, only about half of companies had been offered coverage, and estimated small-business penetration was between 10% and 20%. Henry said small-business growth represents both a market opportunity and a potential source of additional aggregation risk.

Henry said the report had received coverage in 13 articles across 11 outlets and had been downloaded 528 times in approximately two weeks. He said NAIC data provides a market-wide benchmark, while carrier and private datasets generally reflect individual books of business or convenience samples. He emphasized that several findings in the report required external data because the Cyber Supplement does not answer all current supervisory questions.

Schulz said the Cyber Supplement performs its core function by providing premium, policy, and claims information at the market level. However, it provides only partial visibility into admitted and surplus lines results and does not directly report severity by insured size, pricing trends, reinsurance, or net results. He also identified blind spots involving captives, movement of business between offshore and U.S. entities, and cyber-specific expense ratios.

© 2026 National Association of Insurance Commissioners

Schulz and Henry discussed an approved change to Part 5 beginning with the 2027 reporting year. The change will add an admitted/nonadmitted status field for each state or jurisdiction in which an insurer writes business. Henry clarified that the field will indicate whether business is written in a jurisdiction but will not provide state-level financial amounts by admitted or nonadmitted status.

Schulz summarized five broader data gaps: no state-level financial data, no net or reinsurance data, delayed reporting, no quarterly data, and no cyber-specific claim development in Schedule P. Henry added that private claims datasets also have limitations because they may reflect only one insurer's policyholders, omit policy limits, sublimits, and exclusions, and understate costs while claims remain open.

Peterson asked how regulators should interpret increasing claims alongside falling prices and whether the trend could raise solvency concerns. Schulz said cyber insurers remained profitable, although margins may be narrowing. He said increased competition in a growing and maturing market likely contributed to pricing pressure. Peterson and Schulz invited members to identify which data gaps were most relevant to financial analysis, market conduct, and other supervisory work, as well as how states address the timeliness limitations of the current data.

Peterson said the Working Group had additional work to do to determine how the data could support supervision, which gaps were consequential, and what next steps were appropriate. He identified possible targeted regulator discussions, coordination with the Blanks (E) Working Group, and further evaluation of potential reporting changes. No reporting change was proposed or adopted during the meeting.

Having no further business, the Cybersecurity (H) Working Group adjourned.

© 2026 National Association of Insurance Commissioners

Draft: 11/3/26

Cybersecurity (H) Working Group
Virtual Meeting
October 14, 2026

The Cybersecurity (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met Oct. 14, 2026. The following Working Group members participated: Michael Peterson, Chair (VA); Colton Schulz, Vice Chair (ND); Alex Romero and Julia Jette (AK); Chris Erwin (AR); Lori Dreaver Munn (AZ); Damon Diederich (CA); Wanchin Chou, George Bradner, and Kurt Swan (CT); Tim Li (DE); Matt Kilgallen and Paula Shamburger (GA); Daniel Mathis (IA); Nina Hunter (LA); Mary Kwei (MD); Danielle Torres (MI); Gregory Maus and Bubba Aguirre (MN); Troy Smith (MT); Jessica Price (NC); Martin Swanson (NE); Joshua Hilliard (NH); Gille Ann Rabbin (NY); John Pollack and Matt Walsh (OH); Curt Kwak (WA); Jamie Adams and Barbara Belling (WI); and Lela D. Ladd (WY). Also participating were Shannon Hohl (ID); Sandra Darby (ME); Brian Fjeldheim and John Haworth (OR); Matt Gendron (RI); and Tony Dorschner (SD).

1. Heard a Presentation on the Frontier AI Security Readiness and the CrowdStrike Whitepaper

Peterson introduced Brandon Wales (CrowdStrike), noting that Wales leads government affairs and policy engagement across the U.S., Canada, and Latin America and has more than 25 years of national security and cybersecurity experience. Peterson said the presentation would address how artificial intelligence (AI) is changing the speed and economics of cyber risk, shortening the vulnerability life cycle, and creating a new attack surface. Wales presented CrowdStrike's five-step frontier AI security readiness framework.

Wales said AI has not made existing cybersecurity controls obsolete. Instead, it has compressed the time available to identify, assess, and respond to threats. He cited CrowdStrike research showing an 89% year-over-year increase in AI-enabled attacks, a 42% increase in zero-day vulnerabilities exploited before public disclosure, an average e-crime breakout time of 29 minutes, and a fastest observed breakout time of 27 seconds. He asked whether periodic scanning, quarterly assessments, lengthy vulnerability triage, and annual exercises remain appropriate for a machine-speed threat environment.

Wales described a collapsing vulnerability life cycle. He said frontier AI models can analyze large code bases and identify vulnerabilities at a scale that can overwhelm traditional patch-management processes. CrowdStrike identified more than 2,400 potential vulnerabilities across 43 products in one AI-driven research effort. Wales recommended moving beyond vulnerability counts and patching service-level agreements toward continuous identification and reduction of exposures that could cause material harm. He said organizations should connect technical weaknesses to critical systems and business functions and validate whether an adversary could turn an exposure into an incident.

Wales cautioned that frontier AI risk is not solely a vulnerability-management issue. He said adversaries often prefer stolen credentials because legitimate access can be harder to detect and can support follow-on activity. He also described cross-domain attacks, which can exploit fragmented visibility and disconnected security products. Organizations should be able to detect and respond across complex attack life cycles rather than treating individual vulnerabilities as the primary measure of readiness. Wales outlined five elements of the frontier AI security readiness framework:

- **Prioritize by exploitability.** Organizations should focus first on weaknesses that are most likely to be exploited and could produce material harm.

© 2026 National Association of Insurance Commissioners

- **Continuously validate exposure.** Organizations should assess changing network weaknesses on an ongoing basis rather than at a single point in time.
- **Enforce strong identity controls.** Organizations should address human, machine, and AI-agent identities and the trust relationships built into their networks.
- **Detect and respond at machine speed.** Organizations need technology, processes, and personnel capable of responding in real time. Wales said a 27-second breakout cannot be addressed through manual alert triage alone.
- **Apply AI deliberately and securely.** Organizations should secure AI at deployment and maintain visibility into sanctioned and unsanctioned AI, agent behavior, data movement, and malicious activity.

Wales said AI has become an attack surface rather than only a tool. He described risks from malicious code packages, adversaries' use of an organization's AI systems after gaining access to a network, unsanctioned AI, autonomous agent activity, and sensitive data moving into or out of systems. He said governance is necessary but must be supported by technical guardrails. He also characterized AI detection and response technology as available but developing.

Wales compared the current deployment of AI security controls to the federal government's adoption of endpoint detection and response after the SolarWinds incident. He said organizations have deployed AI rapidly without always implementing technology to monitor and secure it. He noted that the National Institute of Standards and Technology (NIST) AI Risk Management Framework is a useful starting point but does not yet have the same body of detailed controls that supports the NIST Cybersecurity Framework.

Wales also discussed AI's defensive benefits. He said cybersecurity tools have incorporated statistical analysis and machine learning for years, helping the industry integrate newer AI capabilities. He identified two promising applications: AI-assisted security operations that triage and act on more alerts and adversarial co-evolution that trains offensive and defensive models together to identify weaknesses and improve defenses. Wales said regulatory programs should recognize defensive AI as an important component of security programs operating at machine speed.

Wales said advanced defensive capabilities should be broadly accessible. Smaller entities may receive these benefits through managed security service providers (MSSPs), managed detection and response (MDR) services, and improvements embedded in underlying technologies. He said the insurance industry can influence cybersecurity practices across many sectors through due diligence, risk analysis, underwriting, and pricing.

Schulz asked what insurance departments should review when insurers use AI coding tools for internal application development. Wales recommended documented governance, monitoring of what tools download and transmit, and supply-chain risk-management practices comparable to those used for external providers. He also recommended using AI to test AI-generated applications for weaknesses and documenting how autonomous tools operate, including the identities and permissions assigned to them. Wales said communities of regulators can help establish common minimum expectations and promote effective practices.

Kilgallen asked whether the NIST AI Risk Management Framework is a sufficient framework to follow. Wales said the framework remains under development and needs more detailed controls to support practical implementation. He said organizations should use it to guide initial risk-management decisions while recognizing that they may need to go beyond it in areas where controls are not yet mature. Debra Decker (Stimson Center) also asked when an update to the framework would be published. Wales said he did not know when the update would be published.

© 2026 National Association of Insurance Commissioners

Hunter asked for minimum expectations that could apply to insurers of different sizes and levels of sophistication. Wales recommended two baseline questions. First, regulators should ask whether an insurer has a documented policy for managing AI and follows it. Second, regulators should ask whether the insurer has deployed technology to monitor and secure the AI it uses. He said organizations cannot credibly claim to use AI securely if they cannot observe its activity, block malicious behavior, and respond to risks.

Wales said small businesses without full-time cybersecurity programs should use an MSSP or MDR service. He said adversaries target internet-accessible weaknesses regardless of an organization's size. Wales also discussed cyber insurers that combine coverage with monitoring, scanning, or MDR services. He said this model can improve security for smaller policyholders and increase insurers' visibility into risk.

Decker discussed whether lenders could encourage or require small and medium-sized business borrowers to obtain cyber insurance. She said cyber insurance underwriting can promote stronger security practices and enable faster adaptation than regulatory requirements. Wales said cyber insurance and managed security services are important risk-management tools for small businesses. He noted that lender requirements could become more likely if cyber incidents produce significant loan losses.

Bradner said Small Business Administration loan programs may provide one avenue for considering cyber insurance requirements. He also identified affordability, availability, and policyholder education as important considerations. He compared the issue to consumer misunderstandings about whether homeowners insurance covers flood losses and said some businesses may similarly assume that commercial insurance includes cyber coverage.

Peterson asked how analysts and examiners should assess both security for AI and AI used for security. Wales said governance and technology should be reviewed together because technology without governance may be directed toward the wrong risks, while governance without technology may not operate at the speed required. He recommended assessing how companies deploy AI to achieve business outcomes, incorporate security into deployment decisions, use AI within security operations, and obtain capabilities through third-party MSSPs or MDR providers. Wales said smaller entities may require more education about the questions they should ask their service providers.

Bradner asked whether the current cybersecurity processes address AI-related risks. Wales said traditional cybersecurity principles and controls remain relevant because AI-enabled attacks still rely on detectable actions, such as credential theft, process execution, and data access. However, organizations may need enhanced capabilities to address increased speed and scale and to secure new attack vectors created by AI deployment.

2. Discussed Other Matters

Peterson said the Working Group's next meeting would include a presentation on the 2025 Cybersecurity Insurance Coverage Supplement data and new ThoughtSpot developments that support regulators' access to annual cyber supplement data. He invited members to submit topics for future Working Group consideration.

Having no further business, the Cybersecurity (H) Working Group adjourned.

© 2026 National Association of Insurance Commissioners

2026 Fall National Meeting
Dallas, Texas

PRIVACY PROTECTIONS (H) WORKING GROUP
Monday, November 16, 2026
4:00 p.m. – 5:00 p.m.

Meeting Summary Report

The Privacy Protections (H) Working Group met Nov. 16, 2026. During this meeting, the Working Group:

1. Adopted its Summer National Meeting minutes (*see NAIC Proceedings – Summer 2026, Innovation, Cybersecurity, and Technology (H) Committee, Attachment Four*).

2. Discussed comments received on the revised full draft of the *Privacy of Consumer Financial and Health Information Regulation* (#672).

2026 Fall National Meeting
Dallas, Texas

THIRD-PARTY DATA AND MODELS (H) WORKING GROUP
Monday, November 16, 2026
8:00 – 9:00 a.m.

Meeting Summary Report

The Third-Party Data and Models (H) Working Group met Nov. 16, 2026. During this meeting, the Working Group:

1. Adopted its Summer National Meeting minutes (*see NAIC Proceedings – Summer 2026, Innovation, Cybersecurity, and Technology (H) Committee, Attachment Five*).

2. Reported that it met Oct. 28 and Sept. 16 in regulator-to-regulator session, pursuant to paragraph 3 (specific companies, entities, or individuals) of the NAIC Policy Statement on Open Meetings. During those meetings, the Working Group took the following action:
   A. Discussed revisions to the property/casualty (P/C) pricing and underwriting regulatory framework and confidentiality.

3. Exposed the P/C pricing and underwriting regulatory framework and confidentiality pathway for a 45-day public comment period ending Dec. 31.

4. Discussed its next steps.

Artificial Intelligence (AI) Risk Evaluation Supplement

Background, Intent, and Instructions

Background:

The rapid expansion of big data and the adoption of artificial intelligence and machine learning are significantly transforming insurance practices. These technologies can offer substantial benefits to both insurance companies and consumers by facilitating the development of innovative products, improving customer interface and enhancing service, simplifying and automating processes, and promoting efficiency and accuracy. However, without robust governance and effective controls, the use of *AI Systems* carries a risk that could lead to Adverse Consumer Impacts or compromise the financial soundness of an insurance company. Insurers are responsible for managing the risks associated with the development and implementation of *AI Systems* and must demonstrate to regulators that appropriate risk-based oversight mechanisms are in place and are functioning effectively. Ultimately, insurers that make decisions or perform actions impacting consumers or financial soundness that are made or supported by advanced analytical and computational technologies, including from *AI Systems*, must comply with all applicable insurance laws and regulations.

Intent:

The NAIC's Innovation, Cybersecurity and Technology (H) Committee charged the Big Data and AI Working Group (BDAIWG) to create resources that would assist regulators to identify and assess both financial and consumer risks from the use of *AI Systems*.

The following exhibits are designed to supplement existing market conduct, financial analysis, and/or financial examination review procedures for reviewing *AI Systems*. As this supplements existing NAIC resources, regulators should continue to consider existing NAIC resources as authoritative but may consider drawing from this Supplement to assist in understanding and assessing a company's use of *AI Systems*. Inquiries and information requests performed related to this Supplement will be coordinated consistent with the guidance provided by the *Market Regulation Handbook (MRH), Financial Condition Examiners Handbook (FCEH)*, and the *Financial Analysis Handbook (FAH)*. Consistent with the Handbooks, this Supplement does not independently create substantive legal duties, standards of conduct, admissions, or expanded regulatory authority.

Moreover, consistent with the coordination practices described in the FCEH, MRH, and FAH, information requests made leveraging the Supplement should be coordinated as much as possible.

With the intention of the document as a supplement acknowledged, there may be some important context that is not specified in this document, particularly related to deciding which companies are in scope of inquiry. This is to avoid the perception of conflicting guidance as many important decisions leading to the full use of this Supplement would be governed by the aforementioned Handbooks.

DRAFT

These optional exhibits allow regulators to determine the extent of *AI Systems* usage for an insurance company and whether additional analysis is needed focusing on financial and consumer risk.

This Supplement consists of the following exhibit questionnaires:

- **Exhibit A: Quantify Regulated Entity's Use of *AI Systems***
- **Exhibit B: AIS Program (Two Options: Narrative or Checklist)**
- **Exhibit C: High-Risk AI Models**
- **Exhibit D: Data Used in High-Risk AI Models**

Regulator Instructions:

Information obtained from the responses to the following Exhibits may supplement guidance and resources used during an existing market conduct, certification, financial analysis, and/or financial examination review, to enhance the regulator's understanding of the extent of *AI Systems* utilization and assessment of potential risk to consumers and the insurance company. Effective assessment requires regulators to maintain a fluent understanding of the applicable laws and regulations, including but not limited to those pertaining to unfair trade practices, unfair claims settlement practices, unfair discrimination protections, corporate governance annual disclosures, confidentiality, privacy protections, financial reporting, and rating.

This Supplement is designed with the intention that regulators using this document will most often use Exhibit A. Based on responses received, regulators may thereafter leverage Exhibit B to inquire about a company's governance practices. Finally, for a subset of a company's use of AI, companies may receive Exhibit C and Exhibit D related inquiries where regulators wish to probe further.

Therefore, regulators using this Supplement may wish to first use Exhibit A to determine if further inquiry using Exhibits B-D is necessary. It is possible that while the company responding indicates usage of *AI Systems*, its use of AI is limited or low in inherent risk such that further inquiry as contemplated by subsequent exhibits is not warranted. For instance, if a regulator is using the Supplement to inquire about a company's claims practices and based on responses to Exhibit A, the regulator determines that the company's use of AI is out of scope of a claims-related targeted exam, the regulator may determine it is appropriate to thereafter rely on the MRH to guide the remaining work.

If information requested through this Supplement has already been provided to this department or any other state department of insurance, the company's response should so state, in which case the regulators may accept prior submissions if the prior response is still current and applicable. Where companies have existing documentation that may satisfy regulatory inquiries, regulators may consider accepting company documentation consistent with guidance described in the FCEH, MRH, and FAH. Where substantially the same *AI Systems* and issues have recently been reviewed, other departments may leverage those submissions and regulatory work consistent with guidance described in the FCEH, MRH, and FAH.

DRAFT

The responses to these exhibits will be considered by regulators when identifying the inherent risks of the insurer. They should also affect the planned examination or inquiry approach, as well as the nature, timing and extent of any further procedures performed.

General Guidance

Machine Learning (ML) Models

Machine learning models are widely used to estimate the relationship between a dependent variable and one or more independent variables. For example, an insurer may want to estimate how a driver's age relates to the probability of being involved in a car accident. There are many ML models that can be used for this purpose.

In particular, generalized linear models (GLMs) are one such technique. As with other machine learning techniques, GLMs are not without risk of causing unfair discrimination or other Adverse Consumer Impacts and require effective governance over data quality, model development, validation, implementation, ongoing performance monitoring, and compliance with applicable laws and regulations. This supplement asks about a company's use of specific models, when utilized within an *AI System*, in Exhibit A to give regulators flexibility in determining their scope of interest.

Based on existing regulatory oversight provided via rate making, regulators may exclude some or all GLMs and other predictive models such as rule-based systems, deterministic algorithms, and fixed mathematical formulas from the Supplement if appropriate. Ultimately, this Supplement is designed with an intention to help regulators understand how models are used with the model type being one aspect in deciding whether further inquiry is necessary.

AI System vs AI Model

Definitions are provided for both terms in the Definitions and Appendix section of the Supplement. Generally, this Supplement is designed with a focus on *AI Systems*. In relation to each other, this Supplement and each of the terms are defined such that an *AI System* is viewed as the primary unit at which governance and reporting discussions will take place; a model is a component of an *AI System*; and a stand-alone model may be examined separately when it independently creates direct consumer impact or material financial risk.

However, there may be times where inquiries are directed at an *AI System* and are better addressed at the *AI Model* level or vice versa. In those instances, companies may engage with department officials to ensure information provided meets expectations and assists in accomplishing regulatory objectives.

Human in the Loop

A core concept in managing AI related risks is related to the degree of involvement of humans in the process. This is often referred to as "human in the loop" when a human is involved in the decision making process or when humans are accountable for a process but not specifically involved in every decision. This practice is more common when agentic AI is deployed.

DRAFT

In evaluating the design and operating effectiveness of human oversight related controls, regulators should consider:

- Authority – Whether the reviewer can reverse, modify, or escalate the AI-assisted decision, and the record shows that overrides actually occur and are documented with reasons.
- Qualifications – Whether the reviewer has the experience or qualifications to evaluate the performance of the AI being used in a process.
- Capacity and Evidence – Whether the design of the human review supports effective oversight and can be demonstrated via override rates, escalation rates, etc.

Reviewing Company Documentation

In several instances, the Supplement suggests a mix of both regulatory inquiries and requests for documentation as a means of allowing regulators to corroborate management's assertions. As regulators decide whether inquiry alone is sufficient or if corroboration is needed, regulators should consult with the MRH, FCEH, and FAH, and any other NAIC resources as needed.

While the Supplement is designed to support regulators, it is not designed as a stand-alone resource and therefore does not provide guidance on standards of fieldwork or related practices.

Materiality

This Supplement was designed to allow regulators the flexibility to either specify materiality or rely on the company's *materiality* calculations as companies respond to Supplement related inquiries.

The concept of *materiality* leverages the definition otherwise used to drive financial exam-related decisions and as defined in the FCEH. For purposes of completing responses related to this Supplement, regulators may opt to rely on a company's internal materiality thresholds provided the information supporting a company's *materiality*, including an explanation for why that threshold was considered appropriate.

However, nothing in the guidance provided within this Supplement should be construed to supersede FCEH guidance.

Risk Assessment

This Supplement also, at times, relies on a company's Risk Assessment. The concept of Risk Assessment centers around the discussion of *Inherent Risk* which leverages the definition reflected in the FCEH. For purposes of completing responses to this Supplement, regulators ask that Risk Assessment information provided to regulators include explicit indication of *Inherent Risk* Assessments (i.e., risk before the consideration of controls or mitigation measures).

Consistent with the discussion of *materiality*, nothing in the guidance provided within this Supplement should be construed to supersede FCEH guidance as it relates to discussions of risk assessment.

Confidentiality

The information requested via these Exhibits will be held confidential, consistent with the requesting state's relevant authority.

Regulators using any of these exhibits should cite examination or other authority as appropriate when requesting information from insurers. Regulators should cite all relevant confidentiality statutes or other specific protections related to documents, materials or other information in the possession or control of regulators that are obtained by or disclosed to the regulators or any other person in the course of a market conduct inquiry and all information reported or provided to the regulator pursuant to cited examination or other authority.

DRAFT

Exhibit A: Quantify Regulated Entity's Use of *AI Systems*

Purpose: To obtain information pertaining to the number of *AI Systems* and *AI Models* used by a company to help facilitate risk assessment. Based on the responses from the company, regulators may ask for additional information related to their *AIS Program* (Exhibit B), high-risk models (Exhibit C), and the types of data (Exhibit D).

Company Instructions:

- Provide the most current counts and use cases of the following, as requested.
- Include all lines of business. If the governance differs by company, line of business, or state, work with your regulator to determine if multiple submissions are needed.
- If a company has an existing *AI System* and/or *AI Model* inventory, the company should discuss with the regulator whether that documentation satisfies the regulator's requests. Regulators will look for columns with information on:
    - *AI System* and/or Model Name
    - Description of how the *AI System* and/or Model is used and the broader *AI System* in which it is contained
    - Operation or Program Area(s) the *AI System* and/or Model is Used Within
    - Inherent Risk Level Based Upon the Company's Risk Assessment
- If a company has questions about any of the fields, please submit them in writing to <regulator contact email address>.

Regulator Instructions:

- Regulators should customize this Exhibit to limit information requested to more targeted inquiries for use in a limited scope exam.
- Depending on the information provided in a company's *AI System* and/or Model inventory, regulators may consider if a separate Exhibit A submission is necessary.
- For column 3, regulators should specify the number of months.

Company Legal Name and Group Name: _____

NAIC CoCode and Group Code: _____

Company Contact Name: _____ Email: _____

Line(s) of Business for Which This Response Applies: _____

DRAFT

Date Form Completed: _____

(1) Use of *AI System* in Operations or Program Area	(2) Number of *AI Systems* Currently in Use	(3) Number of *AI Systems* Implemented in the Past X Months	(4) Use Case *Describe how the systems are used in your business and whether there is a direct consumer impact or a material financial impact*	Do any of the *AI Systems* utilize any of the following (check all that apply)			
				(5) *Generative AI*	(6) *Agentic AI*	(7) Generalized Linear Models (GLMs)	(8) *Other AI/ML Models*
Marketing							
Premium Quotes & Discounts							
Underwriting/ Eligibility							
Ratemaking/ Rate Classification/ Schedule Rating/ Premium Audits							
Claims/ Adjudication*							
Customer Service							
Utilization Management/ Utilization Review/ Prior							

DRAFT

(1) **Use of *AI System* in Operations or Program Area**	(2) **Number of *AI Systems* Currently in Use**	(3) **Number of *AI Systems* Implemented in the Past X Months**	(4) **Use Case** *Describe how the systems are used in your business and whether there is a direct consumer impact or a material financial impact*	**Do any of the *AI Systems* utilize any of the following (check all that apply)**			
				(5) ***Generative AI***	(6) ***Agentic AI***	(7) **Generalized Linear Models (GLMs)**	(8) ***Other AI/ML Models***
Authorization/ Level of Care Determination							
Fraud/Waste & Abuse							
Investment/ Capital Management							
Reserves/ Valuations							
Catastrophe Triage							
Reinsurance							
Other Insurance Practices** (if applicable)							

** Includes Salvage/Subrogation, adjuster assignments, etc.*

*** Any additional operation or program area with the potential to result in either Direct Consumer Impact or Material Financial Impact.*

DRAFT

Exhibit B: (Narrative) AIS Program

Purpose: To obtain the Company's *AIS Program*, including how the company handles governance, risk identification, mitigation, and management framework and internal controls, and third-party *AI Systems* and data for *AI Systems*; and the process for acquiring, using, or relying on third-party *AI Systems* and data. Market and financial regulators should coordinate to gain access to the relevant section of the policies governing the use of *AI Systems*.

Company Instructions: Below are both Checklist and Narrative sections of Exhibit B of the NAIC's AI Risk Evaluation Supplement. If the responses to the narrative indicate that an explanation of a policy, procedure or process can be found on a given page of a document, indicate the page number and the document title and include the document referenced in the submission. If other sections of the document are irrelevant, the relevant sections of the document may be submitted in lieu of the entire document.

For company contact information, please identify an appropriate contact, function, or team that can coordinate responses to follow-up regulatory inquiries and connect regulators with relevant subject matter experts, as appropriate.

Please see definitions at the end of this document.

Regulator Instructions: Regulators should customize this supplement to limit information requested to more targeted inquiries for use in a limited scope exam. The Exhibits are designed so that regulators would use either the Narrative or the Checklist Exhibit but not both.

The references to, and questions about, elements of a company's *AIS Program* are not intended to be interpreted as creating new requirements for *AIS Program*.

Group and Company Legal Name: _____

NAIC Group and Company CoCodes: _____

Company Contact Name: _____ Email: _____

Date Form Completed: _____

1. Provide your *AIS Program*. Click or tap here to enter text.
    a. What role maintains the *AIS Program*? Click or tap here to enter text.

DRAFT

b. Discuss the governance structure, Board reporting and frequency. Click or tap here to enter text.
c. Discuss the process by which the *AIS Program* is integrated throughout the organization, assessed and remediated. Click or tap here to enter text.
d. Discuss the process by which the effectiveness of the *AIS Program, AI Systems,* and individual models are assessed and modified. Click or tap here to enter text.
e. Discuss how responsibility for governance within the organization is assigned and how the organization ensures consistency and alignment. ~~Please include a detailed list of subject matter experts who will be able to answer follow up regulatory inquiries.~~ Click or tap here to enter text.
f. Discuss the integration of the *AI Systems* in the Own Risk and Solvency Assessment (ORSA) and Enterprise Risk Management (ERM) assessments. Click or tap here to enter text.
g. How does the insurance company assess autonomy, reversibility, and reporting impact risk of *AI Systems*? Click or tap here to enter text.
h. Discuss the company's materiality threshold and risk assessment process in evaluating its use of AI. Click or tap here to enter text.

2. Discuss the uses of *AI Systems* that:
   a. Generate a *material financial impact*. Click or tap here to enter text.
   b. Generate a *direct consumer impact*. The company's response should include, but not be limited to, whether there is a human in the loop, whether there are high volume transactions, how the company assures itself that the level of human review is appropriate, and how any exceptions or errors are found and resolved. Click or tap here to enter text.
   c. Generate or impact material information reported in financial statements. Click or tap here to enter text.
   d. Generate or impact risk or control assessment. Click or tap here to enter text.
   e. Discuss the development, testing, and implementation of material *AI Systems* that the Company has implemented. If appropriate, include details regarding where any systems differ from established IT systems and data handling protocols. Discuss the basis for deviation from established practices. Click or tap here to enter text.

3. Discuss how the company:
   a. Assesses, mitigates, and evaluates inherent risks related to the use of *AI Systems* including compliance with existing insurance law. Click or tap here to enter text.
   b. Considers *AI System* risks within the software development lifecycle (SDLC) or within the broader *AIS Program*. Click or tap here to enter text.
   c. Determines whether validation and testing are performed by someone that is independent from development. Click or tap here to enter text.

DRAFT

   d. Identifies material changes that trigger approval, revalidation, or enhanced monitoring. Click or tap here to enter text.

4. Provide the policy for, and discuss the use and oversight of, material *AI System* vendors, model design and testing:
   a. Discuss the validation and testing procedures performed on internally-developed *AI Systems*, including whether validation and testing are performed by someone that is independent from development. Click or tap here to enter text.
   b. Discuss the validation and testing procedures performed on third-party vendor-supplied *AI Systems*. Click or tap here to enter text.
   c. Discuss the testing and verification that has occurred including frequency, scope and methodology. Click or tap here to enter text.
   d. Discuss how the company ensures explainability and transparency into any *AI models* recommending or making decisions. Click or tap here to enter text.

5. Provide the policy for, and discuss the use and oversight of, material *AI Systems* by third-party providers including actuarial, claim, MGA, audit, and/or other professional services. Click or tap here to enter text.
   a. Discuss the testing and verification that has occurred, frequency, scope, and methodology. Click or tap here to enter text.

6. Discuss additional aspects of the framework design and evaluation pertaining to *AI Systems*. Click or tap here to enter text.
   a. Discuss the unit(s) responsible for the *AIS Program*, assessment approach and frequency, and involvement with the program area to the extent it differs from that discussed previously. Click or tap here to enter text.

7. Discuss the criteria and documentation available for identifying a material change and for triggering approval, revalidation, or enhanced monitoring, including pre-deployment testing. Criteria may include, as applicable, periodic or event-triggered reassessment, changes in operational outcomes, etc. Click or tap here to enter text.

8. How does your company:
   a. Discover unsanctioned AI use. Click or tap here to enter text.
   b. Prevent nonpublic personal information from entering unapproved *AI Models*. Click or tap here to enter text.
   c. Govern enterprise assistants that have AI capabilities and support many low- and medium-risk uses. Click or tap here to enter text.
   d. Decide when a general-purpose tool with AI capabilities becomes a high-risk use because of the workflow in which it is deployed. Click or tap here to enter text.

DRAFT

Exhibit B: (Checklist) AIS Program

Purpose: To obtain information about the Company's *AIS Program* including information about how the company handles governance, risk management and internal controls, and third-party *AI Systems* and data for *AI Systems*; and the process for acquiring, using, or relying on third-party *AI Systems* and data. The references to, and questions about, elements of a company's *AI Systems* Governance Risk Assessments are not intended to be interpreted as creating new requirements for *AI Systems* Governance Risk Assessments.

Company Instructions: Provide responses to the questions regarding how the governance of *AI Systems* fits within your company's system of supervision or Enterprise Risk Management program. When referencing documents that are responsive to this exhibit, please include the documents referenced in the submission. Include all companies and lines of business. If the governance differs by entity, line of business, or state, work with your domestic regulator to determine if multiple submissions are needed. See definitions below.

For company contact information, please identify an appropriate contact, function, or team that can coordinate responses to follow-up regulatory inquiries and connect regulators with relevant subject matter experts, as appropriate.

Regulator Instructions: Regulators should customize this supplement to limit information requested to more targeted inquiries for use in a limited scope exam. The Exhibits are designed so that regulators would use either the Narrative or the Checklist Exhibit but not both.

Group and Company Legal Name(s): _____

NAIC Group and Company Code(s): _____

Company Contact Name: _____ Email: _____

Date Form Completed: _____

Ref	***AI Systems* Use Questions for Company**	**Company Response**
1	Has the company adopted a written *AIS Program*? If yes, when was it adopted, what does it govern, who approved the program, and what is the frequency of review for updating?	
2	Was the Board of Directors or management involved in the adoption of an *AIS Program*?	

DRAFT

	2a. What is the role of the Board of Directors or management in the *AIS Program*?		
	2b. How is accountability assigned within the AIS Program and what training has been provided for those charged with oversight of the AIS Program?		
Ref	***AI Systems* Use Questions for Company**		**Company Response**
3	**Reference the processes and procedures of the Company *AIS Program* that addresses the following:**		
	How the Insurance Company…	**Document Name and Page #**	**If not specified in the *AIS Program*, provide details below:**
	3a. Considers *AI System* risks within its Enterprise Risk Management (ERM)		
	3b. Considers *AI System* risks within the Own Risk and Solvency Assessment Report (ORSA), as applicable.		
	3c. Assesses, mitigates, and evaluates *inherent risks* related to the use of *AI System* including the possibility of the company engaging in unfair trade practices		
	3d. Determined "*materiality*" and what threshold was used for assessing *AI Systems* risk and determined the company's risk assessment processes for its use of *AI Systems*.		
	3e. Quantifies *AI System* risk levels		
	3f. Considers *AI System* risk impact on financial reporting		
	3g. Ensures *AI Systems* are compliant with applicable state and federal laws and regulations, including non-discrimination protections		
	3h. Evaluates the risk of *Adverse Consumer Impact*		

DRAFT

	3i. Considers data privacy and protection of consumer data used in *AI Systems*		
	3j. Considers whether *AI Systems* are a contributing factor to consumer complaints and how the company accounts for this in its response to consumer complaints		
	3k. Promotes consumer awareness of the use of *AI Systems* through disclosures, policies, and procedures for consumer notification, as appropriate		
	3l. Trains employees about *AI System* use, ensures human oversight is appropriate, and defines prohibited practices (if any)		
	3m. Provides standards and guidance for procuring and engaging *AI System* vendors		
	3n. Governs, monitors, tests and ensures transparency of *AI Systems* developed by vendors.		
	3o. Considers *AI System* risks within the software development lifecycle (SDLC) or within the broader *AIS Program*		
	3p. Evaluates whether *AI Systems* are suitable for their intended use and should continue to be used as designed		
	3q. Determines whether validation and testing are performed by someone that is independent from development.		
	3r. Identifies material changes that trigger approval, revalidation, or enhanced monitoring.		

DRAFT

Exhibit C: High-Risk AI Models

Purpose: To obtain detailed information on in-production high-risk *AI Models*, such as models making automated decisions that could cause *material financial impact* or direct consumer impact. *AI Model* risk criteria are set by the insurance company. Note that *materiality* should be a consideration in determining the level of risk.

Company Instructions: Fill in the details for each of the *AI Model(s)* requested. If the governance differs by entity, line of business, or state, work with your domestic regulator to determine if multiple submissions are needed. See definitions below. The template below refers to both *AI Systems* and *AI Models* depending on the information being requested. There may be some instances where a company feels information should be provided in relation to the *AI System* and not the *AI Model* or vice-versa. This should be discussed with regulators as part of the submission process to avoid misunderstanding.

Regulator Instructions: Regulators should customize this supplement to limit information requested to more targeted inquiries for use in a limited scope exam. In requesting Exhibit C responses, regulators may wish to request a walkthrough discussion whereby companies provide regulators with an overview of a given *AI Model* or *AI System* which may help regulators understand the use and potential risks of a given use case.

Group and Company Legal Name(s): _____

NAIC Group and Company Code(s): _____

Company Contact Name: _____ Email: _____

Provide who prepared and who reviewed the responses, by function. _____

Date Form Completed: _____

Ref	AI Model Information Requests	Company Response
1	*AI Model* name and version number.	
2	Use cases, output, and purpose(s) of model	
3	Model type used in the *AI System* (*GLM, generative, or other AI/ML algorithm*)	
4	Model Implementation Date	

DRAFT

5	Model development (internal, or third party, or a mix of both – include vendor name if applicable)	
6	Did the Company purchase the *AI model*, was the *AI model* developed specifically for the Company? Is the Company able to modify the *AI model* as necessary?	
7	Model risk(s) (i.e., describe the potential for *adverse consumer impact*) including for models provided by third parties, how the company obtains sufficient evidence to assess and monitor risk and learns of material changes or incidents.	
8	Model limitations (i.e., limitations on *uses cases* or data inputs outside the domain of the training data, or other guardrails to manage model related risk) and how the company determines if limits are being observed and adhered to.	
9	*Does the AI model* automate, augment, or support decision-making?	
10	Discuss testing and validation of model outputs, including but not limited to, testing for model drift, accuracy, unfair trade practices, unfair discrimination, performance degradation, and other potential Adverse Consumer Impacts. Discussion should cover testing and validation practices both prior to and after deployment.	
11	Last date of model testing or validation. Please include the results and corrective actions as a result of the last model testing, if applicable.	
12	Discuss how the model affects the financial statements, risk assessment or controls.	
13	Discuss how the model is reviewed for compliance with applicable state and federal laws, including but not limited to the unfair trade practices act and unfair claims settlement laws.	

DRAFT

14	To the extent permitted by law, discuss if the company has had any actions taken against them for use of this model. Actions may include, but are not limited to, informal agreements, voluntary compliance plans, administrative complaints, ongoing third-party monitoring, cease and desist, remediation, restitution, fines, penalties, investigations, consent orders or other regulatory agency actions.	

DRAFT

Exhibit D: Data Used in High-Risk AI Models

Purpose: To obtain detailed information of the source(s) and type(s) of data used in high-risk *AI Models* to identify risk of *adverse consumer impact*, unfair trade practices, material financial impact, or material financial reporting impact.

Company Instructions: Provide details below for the data used in *AI Model(s)*. If any of the data elements listed are used in the training or test data as part of the development of an *AI Model*, provide information on whether the data element is sourced internally or whether the data element is sourced from a third party, in which case provide the name of the third-party vendor. Leave blank if a data source is not used in the development of *AI model* for the insurance operation. Include all companies and lines of business. If the governance differs by entity, line of business, or state, work with your domestic regulator to determine if multiple submissions are needed. See definitions below.

Regulator Instructions: Regulators should customize this supplement to limit information requested to more targeted inquiries for use in a limited scope exam. Regulators may also wish to customize this supplement for the line of business as some data types may or may not be pertinent based on the responding company's lines of business. Depending on responses received, a regulator may wish to ask for a data dictionary to gain a more complete understanding of a data used for a given model.

Group and Company Legal Name(s): _____

NAIC Group and Company CoCode(s): _____

Company Contact Name: _____ Email: _____

List the *AI System* or *AI models* which rely on the data described in response to this submission:

_____

Describe the Line of Business for Which This Response Applies (complete one for each line of business):

_____

Date Form Completed: _____

DRAFT

Ref	(1) Type of Data Element Used in AI Model	(2) Describe how the AI Model uses this data including how the output is used within operations; or input data is used for training and testing	(3) Internal Data Source	(4) Third Party Data Source / Vendor Name
1	Aerial Imagery			
2	Age, Gender, Ethnicity/Race			
3	Consumer or Other Type of Insurance/Risk Score			
4	Crime Statistics			
5	Criminal Convictions (Exclude Auto-Related Convictions)			
6	Driving Behavior			
7	Education Level (Including school aptitude scores, etc.)			
8	Facial or Body Detection / Recognition / Analysis			
9	Geocoding (including address, city, county, state, ZIP code, lat/long, MSA/CSA, etc.)			
10	Geo-Demographics (including ZIP/county-based demographic characteristics)			
11	Household Composition			
12	Image/video Analysis			
13	Income			
14	Job History			
15	Loss Experience			
16	Medical, including Biometrics, genetic information, pre-existing conditions, diagnostic data, etc.			
17	Natural Catastrophe Hazard (Fire, Wind, Hail, Earthquake, Severe Convective Storms)			

DRAFT

Ref	(1) Type of Data Element Used in AI Model	(2) Describe how the AI Model uses this data including how the output is used within operations; or input data is used for training and testing	(3) Internal Data Source	(4) Third Party Data Source / Vendor Name
18	Online social media, including characteristics for targeted advertising			
19	Personal Financial Information			
20	Reasonable Accommodations or Policy Modifications (granted or requested)			
21	Telematics/Usage-based insurance			
22	Vehicle-Specific Data including VIN characteristics			
23	Voice Analysis			
24	Weather			
25	Other: Non-Traditional Data Elements (Regulator to specify)			

DRAFT

DEFINITIONS AND APPENDIX

Where available, for the purposes of this evaluation terms are defined in accordance with the NAIC Model Bulletin on the Use of *AI Systems* by Insurers (https://content.naic.org/sites/default/files/2023-12-4%252520Model%252520Bulletin_Adopted_0.pdf):

"Adverse Consumer Impact" refers to an *AI System* decision (output) by an insurance company that is subject to insurance law, regulations, and other applicable ~~regulatory~~ standards enforced by the Department that adversely impacts the consumer.

"Agentic AI" refers to *AI Systems* that can pursue objectives across multiple steps without requiring human input at each stage, which could involve perceiving the environment, deciding what to do next, taking actions (often through tools or APIs), observing/analyzing the results, and/or adjusting the action accordingly. The difference is in agency – the system does not just respond to a single prompt or instruction, but actively executes a process.

"Algorithm" means a clearly specified mathematical process for computation; a set of rules that, if followed, will give a prescribed result.

"AI Model" refers to an integral component of an *AI System* that learns from data to perform tasks via computational, statistical and/or machine learning techniques. It combines algorithms, training data and learned parameters to produce outputs from a given set of inputs, such as making predictions or decisions on data it has never seen before.

"AI System" is a machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, content (such as text, images, videos, or sounds), or other output influencing decisions made in real or virtual environments. *AI Systems* are designed to operate with varying levels of autonomy.

"**AIS Program**" refers to the controls and processes that an insurer adopts and implements to ensure the responsible use of *AI Systems*. An *AIS program* may be comprised of general guidelines, governance framework, risk management and internal controls, and third-party oversight.

"Artificial Intelligence (AI)" refers to a branch of computer science that uses data processing systems that perform functions normally associated with human intelligence, such as reasoning, learning, and self-improvement, or the capability of a device to perform functions that are normally associated with human intelligence such as reasoning, learning, and self-improvement. This definition considers machine learning to be a subset of artificial intelligence.

DRAFT

“Augmentation” refers to the use of an *AI System* that suggests an answer and/or advises a human who is making a decision.

“Automation” refers to the use of an *AI System* that does not involve human intervention.

“Direct Consumer Impact” refers to an *AI System* that *augments* or *automates* human decision-making whose output has the potential to create final consumer outcomes including but not limited to underwriting including tier and coverage eligibility determinations, premium determinations, ratemaking, prior authorizations, and claims denial and settlement-related decisions. This has less applicability to *AI Systems* used for internal analytics, model monitoring, and back-office tooling with no consumer-visible output or other internal operational or administrative functions that do not determine or materially affect an outcome for a specific consumer.

“Externally Trained Models” refers to transferred learnings from pre-trained models developed by a third party on external reference datasets.

“Generative Artificial Intelligence (Generative AI)” refers to a class of *AI Systems* that generate content in the form of data, text, images, sounds, or video, that is similar to, but not a direct copy of, pre-existing data or content.

“Generalized Linear Model (GLM)” refers to a statistical modeling technique that extends linear regression by relating predictive variables to an outcome variable using an assumed additive or multiplicative relationship through a link function and an appropriate probability distribution.

“Inherent Risk” refers to the risk of economic loss or inaccurate financial reporting before considering internal controls.

“Internally Trained Models” refers to models developed from data internally obtained by the company.

“Machine Learning (ML)” refers to a field within artificial intelligence that focuses on the ability of computers to learn from provided data without being explicitly programmed.

“**Material Financial Impact**” should be determined based on either the company’s materiality threshold which should be disclosed when responding to Supplement related inquiries or based on the threshold to be provided by regulators.

“**Materiality**” refers to the dollar amount above which the examiner’s perspective of an insurer’s financial position will be influenced.

“Model Drift” refers to the decay of a model’s performance over time arising from underlying changes such as the definitions, distributions, and/or statistical properties between the data used to train the model and the data on which it is deployed.

DRAFT

"Neural Network Model" refers to a machine learning model that mimics the complex functions of the human brain. These types of models consist of interconnected nodes or neurons that process data, learn patterns and enable tasks such as pattern recognition and decision-making, including but not limited to: Single/multi-layer perceptrons/fully connected networks (MLPs/FCs), Deep Learning (DL), Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Long Short-Term Memory Neural Networks (LSTMs), Sequence Models, Large Language Models (LLMs), and Reinforcement Learning Models (RLs).

"Predictive Model" refers to the mining of historic data using algorithms and/or machine learning to identify patterns and predict outcomes that can be used to make or support the making of decisions.

"Residual Risk" refers to an assessment of risk after considering risk-mitigation strategies or controls.

"Support" refers to the use of an *AI System* that provides information but does not suggest a decision or action to a human.

"Third Party" for purposes of this Supplement means an organization other than the insurance company that provides services, data, or other resources related to AI.

"Validation Method" refers to the source of the reference data used for validation, whether Internal, External, or Both.

"Use Case" refers to a description of a specific function in which a product or service is used.

Operations

Claims - Examples: Claim assignment, triage/fast-tracking, individual/bulk claim reserving including loss estimation, imaging/video analysis, fraud detection, litigation, estimation of closure rates, salvage/subrogation, examination/report ordering.

Customer Service - Examples: Agent/broker/internet/customer service interaction (chatbots), online/smart phone apps, loss prevention/risk mitigation advice, payment plans, complaints.

Marketing - Examples: market research, target advertising, market/coverage expansion, customer segment target marketing, demand modeling, agent/broker incentive plans, up/cross-selling.

Other: Cyber Security, Strategic Operations, Reserving, Investments, Capital Management, Financial Reporting, Reinsurance, Legal, Legal Exposure, Reputation Risk.

DRAFT

Ratemaking/Pricing - Examples: Development of overall/base rates, expense/loss loadings, estimation of trends and loss development, development of manual rating factors, tiering criteria, insurance credit scoring, territory boundary definitions, numeric/categorical level groupings and interactions, individual risk rating, telematics/UBI, price optimization, schedule rating factors.

Underwriting - Examples: Policy/coverage acceptance or eligibility, company placement/tiering, schedule rating, decisions based on telematics/UBI, report ordering, retention modeling, inspections, anomaly detection.

DRAFT

PUBLIC COMMENT SUBMISSION

Targeted Redline Recommendations on the NAIC AI Risk Evaluation Supplement — Version 7.0 Draft

TO:	NAIC Big Data and Artificial Intelligence (H) Working Group
VIA EMAIL:	Scott Sobel (ssobel@naic.org) and Miguel Romero (maromero@naic.org)
FROM:	Mebrat Hill, Premium Audit Consultant & Insurance Assurance Practitioner Mebrat.hills@gmail.com | 510-626-7028
DATE:	November 16, 2026
RE:	**Targeted Redlines to AI Risk Evaluation Supplement (Version 7.0 Clean Draft)**

Dear Chair Houdek, Vice Chairs, and Members of the Working Group:

Thank you for the opportunity to comment on the Version 7.0 draft of the Artificial Intelligence (AI) Risk Evaluation Supplement (the "Supplement"). I appreciate the Working Group's continued refinement of the Supplement and its emphasis on risk-based regulatory inquiry, examiner discretion, and coordination with the existing Handbooks.

Several refinements in Version 7.0 are consistent with evidence-based governance principles I outlined in my September 28 and October 28, 2026 submissions, including the importance of evaluating human oversight as an operating control rather than relying solely on the stated presence of a human reviewer.

In particular, the Version 7.0 Human in the Loop guidance is a meaningful refinement. Its focus on reviewer authority, qualifications, and capacity and evidence—including whether overrides and escalations occur and are documented—provides a more practical basis for evaluating the design and operating effectiveness of human-oversight controls.

Two other provisions provide useful foundations that were already present in the Version 6.0 draft and are retained in Version 7.0: the guidance under "Reviewing Company Documentation," which recognizes that regulators may use documentation to corroborate management assertions, and Exhibit B, Narrative Item 2b, which asks how human-review adequacy is assessed in high-volume transactions and how exceptions or errors are found and resolved. My recommendations below build on those existing provisions; I do not characterize them as newly introduced in Version 7.0.

To make these principles more consistently examinable while preserving the Supplement's risk-based and technology-neutral approach, I respectfully submit four targeted, optional redline recommendations. They are intended to clarify the distinction between control design and control operation, support review of selected material decisions, and improve the clarity of relevant data and governing-rule periods. They are not intended to create new substantive legal duties or prescribe a particular technical architecture.

1. General Guidance — "Reviewing Company Documentation" (printed page 4)

Rationale: The existing paragraph recognizes that documentation may be used to corroborate management assertions. The proposed revision clarifies the distinction between evidence that a control was designed and evidence that it operated when required.

Proposed replacement for the first paragraph:

> *In several instances, the Supplement suggests a mix of both regulatory inquiries and requests for documentation as a means of allowing regulators to corroborate management's assertions. Where appropriate to the risk and materiality of the use case, regulators may distinguish between evidence of control design (such as policies, procedures, and assigned responsibilities) and evidence of control operation (such as contemporaneous records and exception handling demonstrating that controls functioned as intended). As regulators decide whether inquiry alone is sufficient or corroboration is needed, regulators should consult with the MRH, FCEH, and FAH, and any other NAIC resources as needed.*

2. Exhibit B — AIS Program: Narrative Item 1d (printed page 10) and Checklist Item 3p (printed page 14)

Rationale: These complementary changes would make the narrative and checklist options more explicit about control exceptions, unauthorized or bypassed actions, and whether material remediation was effective—not merely whether routine assessments occur.

Proposed revised Narrative Item 1d:

> *1d. Discuss the process by which the effectiveness of the AIS Program, AI Systems, and individual models are assessed and modified, including how control exceptions, bypassed reviews, and required remediations are detected, escalated, addressed, and verified.*

Proposed revised Checklist Item 3p:

> *3p. Evaluates whether AI Systems are suitable for their intended use and should continue to be used as designed, including whether controls operate effectively and material exceptions are appropriately resolved.*

3. Exhibit C — High-Risk AI Models: Item 9 (printed page 16)

Rationale: This addition would connect the model-level inquiry to the operational context in which a material decision is made, helping an examiner understand how required human oversight is confirmed and how material decisions are documented.

Proposed revised Item 9:

> *9. Does the AI model automate, augment, or support decision-making? For material decisions, describe how the insurer verifies that required human review or approval occurs before a decision takes effect, and what records document the relevant model version, governing rules, human review, and final disposition.*

4. Exhibit D — Data Used in High-Risk AI Models: Column (2) (printed page 19)

Rationale: For material insurance decisions, the applicable policy period and effective date of governing rules may be as important as the data itself. Identifying relevant periods and versions can reduce ambiguity when an examiner later reviews the decision.

Proposed revised Column (2) heading:

> *(2) Describe how the AI Model uses this data, including how the output is used within operations or how input data is used for training and testing. Where relevant to the use case, identify applicable policy*

periods and effective dates of governing rules, relevant data and rule versions, and how the data relates to resulting decisions.

These focused additions are intended to support consistent, reproducible examination while preserving examiner discretion, existing Handbook authorities, proportionality, and technology neutrality. They rely on evidence appropriate to the risk and materiality of the use case and do not require a prescribed platform, reporting architecture, or universal record format.

The Working Group may use these policy and drafting recommendations; this submission does not disclose, dedicate to the public, or grant any license or rights in any proprietary implementation technology, software, or assurance methodology.

I appreciate the Working Group's consideration and would be pleased to assist with any drafting questions.

Respectfully submitted,

Mebrat Hill
Premium Audit Consultant & Insurance Assurance Practitioner
Mebrat.hills@gmail.com | 510-626-7028

VERDICT
Verdict Systems Inc. · Houston, Texas · founder@verdict.systems · +1 832 641 1892

November 12, 2026

Commissioner Nathan Houdek, Chair
Doug Ommen and Mary Block, Co-Vice Chairs
Big Data and Artificial Intelligence (H) Working Group
National Association of Insurance Commissioners
By email to Scott Sobel (ssobel@naic.org) and Miguel Romero (maromero@naic.org)

Re: Version 7.0 of the AI Risk Evaluation Supplement, and the record behind a human override

Dear Chair Houdek and Members of the Working Group:

Verdict Systems Inc. supports adoption of Version 7.0 of the AI Risk Evaluation Supplement at the Fall National Meeting. This letter asks for no change to the text. The comment period on Version 6.0 closed on October 30, and the draft minutes of the November 2 meeting record that no further substantive changes will be made before adoption.[1] We write about the work those minutes say comes next: bringing the Supplement into the examination handbooks, and further study of agentic AI.

Verdict builds technology that records automated and AI-influenced actions at the moment they occur, in a form a third party can check. We have a commercial interest in this subject and say so at the outset. Nothing below requires a particular vendor or product. We commented on Version 5.0 on September 29 and submitted statements for the record on October 8.

1 What Version 7.0 added

Version 7.0 adds guidance on human review under the heading "Human in the Loop." It tells regulators who are evaluating human oversight to consider whether the reviewer "can reverse, modify, or escalate the AI-assisted decision, and the record shows that overrides actually occur and are documented with reasons." It also asks whether effective oversight "can be demonstrated via override rates, escalation rates, etc."

We support this guidance. A written policy can say that a person reviews the AI system's work. The new language has the examiner look for proof that the person did, case by case. Commenters at the November 2 meeting drew the same line between a control that is documented and evidence that the control operated.[2]

Version 7.0 also says where the next question belongs. It states that the Supplement "does not provide guidance on standards of fieldwork or related practices," and the staff note beside that sentence explains that this is why commenters' input on standards of evidence was not incorporated.[3] We agree with that placement. How an examiner tests a record is fieldwork, and the handbooks are where fieldwork guidance belongs.

2 The question the new guidance raises

"The record shows" assumes a record the examiner can rely on. Examiners already judge this for decisions people make. An adjuster's note entered on the day a claim was denied and a memo written the week the examination was announced can describe the same decision in the same words. An examiner would give more weight to the note, because it was made at the time and before anyone knew it would be reviewed.

With AI systems the two are harder to tell apart. The usual record of what an AI system did, and of what a person did about it, is a log: entries in a file or database inside the company's own systems.

Someone with administrator access can change or delete an entry, and unless the system was built to prevent that, the changed entry looks the same as an original. An override reason typed in after a complaint arrives can carry the same date as one entered when the decision was made.

The same applies to the rates the guidance mentions. An override rate is a count of override entries divided by a count of decisions. If entries can be added or removed afterward, the rate can move without anyone's conduct having changed.

Order of events is the other difficulty, and it grows with agentic AI, which Version 7.0 defines as systems that "pursue objectives across multiple steps without requiring human input at each stage." When a system takes several steps on its own, the examiner needs to know which came first. Did the person approve the payment before it went out, or after? Version 7.0 keeps the definition and the Exhibit A column for agentic AI, and leaves the more detailed agentic questions that appeared in Version 6.0 for further study.[4]

3 Three questions an examiner can ask of any record

An examiner can test a record without a technical background. Three questions cover most of it, and each has an everyday counterpart.

(a) *Was it made at the time?* A record created when the action happened is worth more than one assembled later. The everyday counterpart is a time-stamped receipt set against a note written from memory.

(b) *If someone changed it afterward, would the change show?* Software can compute a short code from a record, called a hash and often described as a digital fingerprint. Change one character of the record and the fingerprint changes. If the fingerprint is kept where the record's owner cannot alter it, any later edit to the record can be detected. A bound ledger with numbered pages works the same way, because a missing or replaced page is obvious.

(c) *Can someone outside the company confirm the first two answers?* Recording a deed with the county is the everyday counterpart. The recorder's office keeps a dated public entry that neither buyer nor seller can go back and change, and anyone can look it up. Public transparency logs do this job for digital records. What goes to the log is the fingerprint, signed by the sender, and the log publishes it with the time it arrived. The record itself, and any consumer information in it, stays with the company. An examiner who is later handed the record can compute its fingerprint and compare it with the public entry. If the two match, the record has not changed since the time shown in the log.

A company can meet all three with openly published methods, and a regulator does not have to choose among them.[5]

4 What such a record does not show

These records have limits. A record that passes all three questions shows that it existed in a given form at a given time and, where records are linked, in a given order. It does not show that the event happened as described, that every action was recorded, or that the decision was correct. Examiners would still test completeness and accuracy by sampling, as they do now. They would be sampling records whose time, integrity and order can be checked without relying on the company's word. A company that answers that it cannot produce records like these has still told the examiner something useful about inherent risk.

5 Where this could fit in next year's work

We offer two suggestions for the work that follows adoption. Neither changes Version 7.0.

Handbook guidance. When discussion resumes on incorporating the Supplement into the Market Regulation Handbook or the Financial Condition Examiners Handbook, the three questions could be written as a short procedure for testing a sampled override. One possible wording:

> SAMPLE WORDING
>
> "For a sample of AI-assisted decisions with Direct Consumer Impact, including decisions the company reports as overridden, obtain the company's record of the AI System's output and of the human review. For each record, determine (a) whether it was created when the action occurred, (b) whether a later change to it could be detected, and (c) whether its time of creation can be confirmed by a party other than the company, without access to the company's systems."

Further study of agentic AI. For systems that act in several steps, a fourth question helps: can the records show the order of the steps? One open method is for each later record to include the fingerprint of the record before it, which fixes the order. Appendix A shows two records linked this way. Section 5.3 of our September 29 letter proposed a definition, "Contemporaneous Action Record," that could serve as shorthand in that study.

6 A worked example staff can check

Appendix A reproduces two linked records of a fictional two-step decision: an AI recommendation, then a human override with a reason. The first check needs only a web browser. The others use standard tools and a public log that Verdict does not operate. The record in Appendix A of our September 29 letter (public log index 2951501338) can still be checked the same way today.

The records were produced on Verdict's public sandbox, which signs with a short-lived key generated for each request. They demonstrate the checks described above. They do not demonstrate the identity or key custody of a production signer.

We ask only that this letter be kept with the Working Group's record for the work ahead. We would welcome questions from members or staff, and would be glad to walk through the example with anyone who wants to try it.

Respectfully submitted,

Shayne Beavan
Founder, Verdict Systems Inc.
founder@verdict.systems · +1 832 641 1892

Disclosure. Verdict Systems Inc. is a Delaware corporation based in Houston, Texas. It has filed U.S. patent applications relating to its products; none has been examined. It is not an insurer or producer. The suggestions above do not depend on any Verdict technology.

1 Draft minutes of the Working Group's November 2, 2026 meeting (draft dated November 10), Attachment A to the materials for the November 16 session. The draft minutes also say the Supplement "will continue to evolve, including through work with the Market Regulation and Consumer Affairs (D) Committee and the Financial Condition (E) Committee, as well as further study of agentic AI," and that discussion of incorporating it into the two handbooks "will continue next year."

2 Same draft minutes, item 3.

3 Comment SS6 in "AI Risk Evaluation Supplement v7.0 (Edited)," posted on the Working Group's page and retrieved November 12, 2026.

4 Version 6.0's Exhibit A question on actions taken "without human approval before the action" and its Exhibit C Ref 3a do not appear in "AI Risk Evaluation Supplement v7.0 (Clean)," retrieved November 12, 2026.

5 Cryptographic hashing, trusted timestamps under RFC 3161, and public append-only transparency logs of the kind web browsers rely on through Certificate Transparency (RFC 6962).

Appendix A. Worked example: an AI recommendation and a human override

DEMONSTRATION · FICTIONAL SCENARIO · PUBLIC SANDBOX TIER

Scenario (fictional). An AI credit model at a fictional bank re-scores a pre-approved loan offer after the customer's salary deposits drop, and recommends withdrawing the offer. Bank policy sends the recommendation to a lending officer before any message reaches the customer. The officer confirms with the employer that salary has resumed and approves a reduced amount with a 60-day review, and the reason is recorded. In the Supplement's words, this is an override "documented with reasons." Each step is its own record, and the second record contains the first record's fingerprint. The example is a loan because the records were built for an earlier demonstration. The mechanism is the same for a claim payment or an underwriting referral.

STEP 1 RECORD	`ser_a383778f88ae4f060279d8ba137926f8`
STEP 1 FINGERPRINT	`65723a7438b7c36686ed7658f8016b2dd5380eaa3500a40a10423fc6a5c9b99c`
STEP 1 PUBLIC LOG	Sigstore Rekor, log index 3247822463 · logged 2026-10-14 02:19:05 UTC · log value `7e0c76c0c9382c5237c3e4322e670623ff51e1b81f7feaab2eca9dc745a54104`
STEP 2 RECORD	`ser_5086c93c22af206a567fa06cf634682e`
STEP 2 FINGERPRINT	`192eb3dfe9fa65765ee508f0e29d00975cbc3370f81419ceef616179ef220ddf`
STEP 2 PUBLIC LOG	Sigstore Rekor, log index 3247822472 · logged 2026-10-14 02:19:06 UTC · log value `8e22075588e58e4111b7dc246a381f27762bfc3721e6fcaf81340ed72d7e6862`

How to check them

1. **In a web browser, with nothing to install.** Go to search.sigstore.dev, choose Log Index as the attribute, and enter 3247822463, then 3247822472. Each page shows the entry's hash, which matches the log value above, and the time the log received it, shown in your own time zone. Sigstore's Rekor is a public transparency log that Verdict does not operate.
2. **On any computer.** The log value is the SHA-256 fingerprint of the record's own fingerprint. Computing it yourself gives the same value, which shows the record's fingerprint existed no later than the log time.

   ```
   printf '%s' 65723a7438b7c36686ed7658f8016b2dd5380eaa3500a40a10423fc6a5c9b99c | shasum -a 256
   printf '%s' 192eb3dfe9fa65765ee508f0e29d00975cbc3370f81419ceef616179ef220ddf | shasum -a 256
   ```

3. **With the full records.** Both records are public. Verdict's open-source verifier recomputes each fingerprint from the record's content and checks it against the log, including the log's own proof that the entry is included, using a Rekor key taken from Sigstore's repository. A change to any byte of a record produces a different fingerprint that no longer matches.

   ```
   curl -sO https://verdict.systems/samples/ser_a383778f88ae4f060279d8ba137926f8.json
   curl -sO https://verdict.systems/samples/ser_5086c93c22af206a567fa06cf634682e.json
   git clone https://github.com/verdictsystems/ser-spec
   curl -sL -o rekor.pem \
     https://raw.githubusercontent.com/sigstore/root-signing/main/targets/rekor.pub
   python3 ser-spec/tools/verify.py ser_a383778f88ae4f060279d8ba137926f8.json \
     --anchor --rekor-key rekor.pem
   python3 ser-spec/tools/verify.py ser_5086c93c22af206a567fa06cf634682e.json \
     --anchor --rekor-key rekor.pem
   ```

   Each run ends with "RESULT: content self-consistent and anchor confirmed."

4. **Order.** The second record's content includes the first record's fingerprint (field `prior_merkle_root`). Because that value is part of the content that produced the second fingerprint, the override record was created after the recommendation it responds to, and neither record can be altered without breaking the match. The check prints True:

```
python3 -c "import json
a = json.load(open('ser_a383778f88ae4f060279d8ba137926f8.json'))
b = json.load(open('ser_5086c93c22af206a567fa06cf634682e.json'))
print(b['record']['metadata']['prior_merkle_root'] == a['merkle_root'])"
```

What these records show, and what they do not

STATEMENT	STATUS
The content of each record is the content that was logged	**Shown**
Each record existed no later than its log time	**Shown**
The override record was created after the recommendation record	**Shown**
The checks run on the examiner's own computer, against a log Verdict does not operate	**Shown**
The loan, the review and the override occurred	**Not shown (fictional)**
Every action in the workflow was recorded	**Not shown**
Who signed the records	**Not shown (sandbox key)**

The last three rows are why an examination would still sample and test.

Draft 10/26/2026
Adopted by the Executive (EX) Committee and Plenary, Pending, 2026
Adopted by the Innovation, Cybersecurity, and Technology (H) Committee, Pending, 2026

2026 Proposed Charges

INNOVATION, CYBERSECURITY, AND TECHNOLOGY (H) COMMITTEE

The mission of the Innovation, Cybersecurity, and Technology (H) Committee is to: 1) provide a forum for state insurance regulators to learn about and have discussions regarding: cybersecurity, innovation, data security and privacy protections, and emerging technology issues; 2) monitor developments in these areas that affect the state insurance regulatory framework; 3) maintain an understanding of evolving practices and use of innovation technologies by insurers and producers in respective lines of business; 4) coordinate NAIC efforts regarding innovation, cybersecurity and privacy, and technology across other committees; and 5) make recommendations and develop regulatory, statutory or guidance updates, as appropriate.

Ongoing Support of NAIC Programs, Products, or Services

The **Innovation, Cybersecurity, and Technology (H) Committee** will:

1. Provide forums, resources and materials related to developments and emerging issues in innovation, cybersecurity, data privacy, and the uses of technology in the insurance industry in order to educate state insurance regulators on these developments and how they affect consumer protection, insurer and producer oversight, marketplace dynamics, and the state-based insurance regulatory framework.
2. Consider and coordinate the development of regulatory guidance and examination standards related to innovation, cybersecurity, data privacy, and the use of big data and artificial intelligence (AI) including machine learning (ML), in the business of insurance, and technology. This includes drafting and revising model laws, white papers, and other recommendations as appropriate.
3. Oversee the work of the Data Call Study Group to study the enhancement of regulator access to high-quality and timely data allowing for evidence-informed decisions, enhanced supervisory capabilities, and improved efficiency.
4. Track the implementation of and issues related to all model laws pertaining to innovation, technology, data privacy, and cybersecurity, including the *Insurance Data Security Model Law* (#668), the *NAIC Insurance Information and Privacy Protection Model Act* (#670), the *Privacy of Consumer Financial and Health Information Regulation* (#672), and the *Unfair Trade Practices Act* (#880).
5. Coordinate and facilitate collaboration with and among other NAIC committees and task forces to promote consistency and efficiency in the development of regulatory policy, education, training, and enforcement materials and tools.
6. Follow the work of federal, state, and international governmental bodies to avoid conflicting standards and practices.

The **Big Data and Artificial Intelligence (H) Working Group** will:

1. Research the use of big data and AI (including ML) in the business of insurance.

2. Monitor state, federal, and international activities on AI, including updates to the regulatory framework for the oversight of AI, to: i) respond to such activities, where appropriate, ii) address potential impacts on existing state insurance laws or regulations, and iii) support adoption of the *Model Bulletin on the Use of AI Systems by Insurers*.

© 2026 National Association of Insurance Commissioners

3. Facilitate discussion on AI systems evaluation, including coordination across the NAIC, to identify and/or develop tools, resources, materials, training, and guidance to assist regulators in their review of AI systems used by licensees.

4. Facilitate and coordinate foundational and contextual educational content for regulators on topics related to the use of big data and AI techniques, tools, and systems in the insurance industry.

The **Cybersecurity (H) Working Group** will:

Cybersecurity Charges

~~1. Monitor cybersecurity trends, such as vulnerabilities, risk management, governance practices, and breaches, that have the potential to affect the insurance industry.~~
~~2. Facilitate communication across state insurance departments regarding cybersecurity risks and events.~~
~~3. Develop and maintain regulatory cybersecurity response guidance to assist state insurance regulators in the investigation of national insurance cyber events.~~
~~4. Monitor federal and international activities on cybersecurity, engaging in efforts to manage and evaluate cybersecurity risk.~~
~~5. Coordinate NAIC committee cybersecurity work, including cybersecurity guidance developed by the Market Conduct Examination Guidelines (D) Working Group and the Information Technology (IT) Examination (E) Working Group.~~
~~6. Work with the Center for Insurance Policy and Research (CIPR) to receive updates on cybersecurity research efforts, by the CIPR and others, and to analyze publicly available cybersecurity-related information.~~
~~7. Support the states with implementation efforts related to the adoption of Model #668 including supporting the research into and, if approved, the implementation of a cybersecurity event notification portal.~~

1. ~~Coordinate with committee support to support cybersecurity preparedness for state insurance regulators, including advising on training development and facilitating tabletop exercises.~~ Threat Intelligence — Monitor cybersecurity trends, vulnerabilities, and risk management-governance practices, that have the potential to affect the insurance industry and share relevant threat and trend information with state insurance departments on an ongoing basis. Refer emerging trends or thematic concerns to the Market Regulation and Consumer Affairs (D) Committee and the Financial Condition (E) Committee, as appropriate.

2. Event Response —
Support the development and maintenance of the NAIC's cybersecurity event response resources and, if approved, carry out the evaluation and coordination role contemplated for the Working Group under the Cybersecurity Event Response Coordination Framework. This includes evaluating cybersecurity events against the Framework's criteria, recommending to (D), (E), and (H) Committee leadership whether an event warrants a coordinated response, and assisting NAIC staff in identifying a coordinating lead state and facilitating information sharing among affected states during the response.

3. Supervisory Integration —
Coordinate NAIC committee cybersecurity work, including cybersecurity guidance developed by the Market Conduct Examination Guidelines (D) Working Group and the Information Technology (IT) Examination (E) Working Group. Strengthen the integration of cybersecurity into insurer supervision by developing mechanisms to incorporate cybersecurity examination findings, Insurance Data Security Model Law (#668) certifications, and cybersecurity event notification data received under applicable state law into ongoing regulatory analysis, and developing baseline operational resilience guidance for insurers.

Commented [MR1]: NAMIC edit adds "if approved" to this wording.

Commented [MR2]: Edit to address NAMIC's inquiry adding "received under applicable state law" into the language.

© 2026 National Association of Insurance Commissioners

4. Model #668 & Notification Portal —
   Support the states with implementation efforts related to the adoption of Model #668, including supporting the research into and, if approved, the implementation of a cybersecurity event notification portal. Such support should include developing standardized data elements for event notifications and identifying any implementation gaps or potential enhancements to Model #668 for referral to the appropriate committee.

~~8.~~5. Regulator Preparedness —
   Coordinate with committee support to support cybersecurity preparedness for state insurance regulators, including advising on training development and facilitating tabletop exercises. Exercises should rehearse the coordination procedures contemplated by the Cybersecurity Event Response Coordination Framework.

Cybersecurity Insurance Charges

1. Market Monitoring —
   Monitor industry trends pertaining to cyber insurance, including meeting with subject matter experts (SMEs) and evaluating data needs of state insurance regulators. ~~Considerations should include the availability and affordability/pricing of cyber insurance, disclosures, limits and sublimits, as well asd sublimits in policies, policy language and trends in requirements, underwriting practices, and the role of reinsurance in the cyber insurance market~~Coordinate with NAIC work groups addressing cyber insurance related issues, such as the Casualty Actuarial and Statistical (C) Task Force.

2. Data Analysis & Reporting —
   Coordinate with committee support to conduct analysis pursuant to the NAIC's Cyber Insurance Report, and develop a guide for states on cyber insurance data analysis best practices.

~~1.~~3. Coverage Supplement & Capital Adequacy
   — Review the NAIC's Property & Casualty Annual Statement Cybersecurity Insurance Coverage Supplement recommending changes and/or developing reports to supplement data development as necessary, in coordination with the Capital Adequacy (E) Task Force.

~~2. Coordinate with NAIC work groups addressing cyber insurance related issues, such as the Casualty Actuarial and Statistical (C) Task Force.~~
~~3. Monitor federal and international activities related to cyber insurance and financing mechanisms for cyber risk.~~
~~4. Coordinate with committee support to conduct analysis pursuant to the NAIC's *Cyber Insurance Report*.~~
~~5. Review the NAIC's *Property & Casualty Annual Statement Cybersecurity Insurance Coverage Supplement* recommending changes and/or developing reports to supplement data development as necessary.~~
~~6. Consider and develop a guide for states on cyber insurance data analysis best practices.~~

The **Privacy Protections (H) Working Group** will:

1. Use state insurance privacy protections regarding the collection, data ownership and use rights, and disclosure of information gathered in connection with insurance transactions to draft a new/revised privacy protections model act to replace/update NAIC models such as Model #670 and/or Model #672.
2. Monitor state, federal, and international activities on privacy, engaging in efforts to manage and evaluate privacy.

The **SupTech/GovTech (H) Subgroup** will:

© 2026 National Association of Insurance Commissioners

1. Facilitate technology, innovation, and SupTech/GovTech presentations from leading technology companies for state insurance regulators to provide them with insights into cutting-edge technology and innovation.
2. Facilitate technology, innovation, and SupTech/GovTech presentations from specialized vendors for state insurance regulators to assist in identifying vendor solutions that may benefit regulators.

The **Third-Party Data and Models (H) Working Group** will:

1. Develop and propose a framework for the regulatory oversight of third-party data and predictive models.
2. Monitor and report on state, federal, and international activities related to governmental oversight and regulation of third-party data and model vendors and their products and services.

NAIC Support Staff: Miguel Romero, Scott Morris, and Scott Sobel

© 2026 National Association of Insurance Commissioners

Hi Miguel,

Thank you for sharing the draft 2027 charges. You can tell a lot of work went into them. The Cybersecurity (H) Working Group's consolidation from fourteen charges to eight reads really well, and the summary memo makes the mapping easy to follow.

One small thought. I noticed the Supervisory Integration charge for the Cybersecurity (H) Working Group doesn't yet touch on AI controls. That area is moving quickly, and by the end of 2027 the landscape could look very different than it does today, so it may be worth keeping on the radar.

If it's helpful, one short phrase in the charge's second sentence would cover it. With the new words added, the sentence would read:

"Strengthen the integration of cybersecurity into insurer supervision by developing mechanisms to incorporate cybersecurity examination findings, Insurance Data Security Model Law (#668) certifications, and cybersecurity event notification data into ongoing regulatory analysis, including findings on controls over consequential actions initiated by automated systems and AI agents, and developing baseline operational resilience guidance for insurers."

The only new words are "including findings on controls over consequential actions initiated by automated systems and AI agents." Nothing else changes.

Thanks again to you and the team for all the good work. Happy to help any way I can.

Tim Fox

President and Founder

EXECUTION RIGHTS

tim@executionrights.com | **executionrights.com**

The authorization layer for consequential AI agent actions.

NAMIC®
NATIONAL ASSOCIATION OF MUTUAL INSURANCE COMPANIES

317.875.5250 | [F] 317.879.8408
10 W. 106th St., Suite 100 | Indianapolis, Indiana 46290
202.628.1558 | [F] 202.628.1601
20 F Street N.W., Suite 510 | Washington, D.C. 20001

November 6, 2026

Commissioner Yaworsky (FL), Chair
NAIC Innovation, Cybersecurity, and Technology (H) Committee
c/o Miguel Romero, Director of Innovation, Cybersecurity and Technology
via email maromero@naic.org

Re: (H) Committee Proposed Charges for 2027

Chair Yaworsky and Members of the Committee:

On behalf of the National Association of Mutual Insurance Companies (NAMIC)[1], we are writing in response to the exposure of the Innovation, Cybersecurity, and Technology (H) Committee proposed charges for 2027. We recognize the intent that the changes to the charges are for clarification purposes; our comments are provided with that intent in mind.

[1] Clarify the Cybersecurity Event Response Coordination Framework is Not Yet Adopted
In Section 2, Event Response, the charge states that the Cybersecurity (H) Working Group will carry out the evaluation and coordination role contemplated under the Cybersecurity Event Response Coordination Framework, with additional details on the activities those roles will entail. NAMIC has been separately engaging with the Market Conduct Examination Guidelines (D) Working Group on the Framework effort, and we note that the Framework has not yet been finalized and adopted. As such, we recommend that the Cybersecurity Charge 2 "Event Response" be edited to clarify the Framework's status. This could be achieved through minor edits as follows:

> Support the development and maintenance of the NAIC's cybersecurity event response resources and, if approved, carry out the evaluation and coordination role contemplated for the Working Group under the Cybersecurity Event Response Coordination Framework. ~~This includes~~ Such support will depend on the final adopted Framework, but might include evaluating cybersecurity events against the Framework's criteria, recommending to (D), (E), and (H) Committee leadership whether an event warrants a coordinated response, and assisting NAIC staff in identifying a coordinating lead state and facilitating information sharing among affected states during the response.

[1] The National Association of Mutual Insurance Companies is the foremost trade association representing the property/casualty insurance industry. Serving more than 1,300 member companies, including local and regional insurers as well as some of the nation's largest carriers, NAMIC members collectively write $453 billion in annual premiums, representing 58 percent of the homeowners and 52 percent of the automobile insurance markets. For more than 130 years, NAMIC has been the leading voice advancing public policy solutions and regulatory frameworks that promote a strong, competitive market and protect member companies and their policyholders.

[2] Clarification Request on the Supervisory Integration Charge
The Supervisory Integration charge appears to incorporate cybersecurity examination findings, Model # 668 certifications, and event notification data into ongoing regulatory analysis, while also calling for baseline operational resilience guidance for insurers. With respect to the ongoing regulatory analysis aspect using cybersecurity event notification data, NAMIC respectfully requests clarification on what the Committee intends by this and where that data might be coming from. Through our participation on the cybersecurity event notification portal effort, our understanding is that if the portal operates as intended, only states that have an event report sent to them should have access to that data, and the broader NAIC would not. Discussion should occur over what limits would govern secondary uses of examination findings, certifications, and incident notifications, including opportunities for insurers to correct or contextualize information. Given the forgoing, NAMIC would appreciate additional insight into where the cybersecurity event notification data would be coming from for purposes of an ongoing regulatory analysis.

Separately, with respect to the baseline operational resilience guidance reference, NAMIC also respectfully requests clarification on what the Committee intends with this aspect of the charge. For instance, would the operational resilience guidance remain advisory in nature only? If it is intended to be anything more than advisory in nature, and become an examination expectation, NAMIC strongly encourages clarification that work would go through appropriate exposure for comment and collaboration with interested parties.

Sincerely,



Lindsey Stephani
Senior Policy Vice President
Data Science, Artificial Intelligence, and Cybersecurity
NAMIC

The AI Landscape: Where the Technology Is Today

Agentic AI has naturally evolved from GenAI LLMs (Large Language Models)

Many clients today

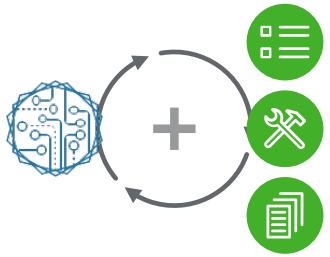
Solo LLM
Raw Generation



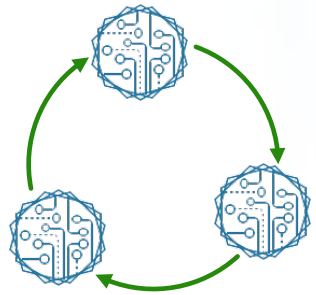
RAG (Retrieval-Augmented Generation)
Reduce Hallucinations



Simple Agents
Self Correct Errors



Multi- Agent Fleets
Breakdown & Solve Complex Use Cases



Limited systems that solve tasks

Enterprise adoption

2022 | **2023** | **2024** | **2025+**

Intelligent solutions that solve full workflows

AI generates content and can **decide the output** based on input guidance e.g., ideas/code, translate, summarize, format

AI generates with **contextually relevant guidance** retrieved from **memory**

AI can **plan, use tools** and **memory**, and iteratively **reflect** to act and perform multi-step **reasoning**

Multiple Agents work **collaboratively** to solve complex tasks. Guardian agents can take on reviewer and monitoring roles.

Insurance-Relevant Examples	*Drafts a policyholder letter or claim summary from a prompt - adjuster reviews before it is sent*	*Underwriters use RAG to answer coverage questions against policy forms, filings and underwriting guidelines*	*Triages a claim - checks coverage, flags fraud indicators, and routes unclear cases to an adjuster*	*Agents run claims intake: one validates coverage, another scores fraud risk, a guardian agent checks decisions for bias*

Copyright © 2026 Deloitte Development LLC. All rights reserved.

Incorporating technology into the audit

Here are some of the ways technology, including AI, is used at Deloitte:

Programmatic data acquisition & reconciliation
Reduced manual effort and fewer errors through automated mapping of your data to our systems

Raw data to common data model
Less work for your team as we translate your data into our audit-ready format automatically

Accelerated data visualization
Faster insights into your financial trends and metrics

Walkthroughs to draft documentation
Walkthrough notes can more efficiently be converted to draft documentation with the potential to reduce repeated follow-ups

On-demand pre-review
Fewer review cycles and actionable suggestions and feedback, reducing back-and-forth during the audit

Draft accounting memos
Earlier visibility into key accounting complexities, helping reduce last-minute surprises

Draft reports and representation letters
Cleaner, more complete draft documents and reports saving your team time during year-end



Analyst report sensing
Sensing and benchmarking analyst reports to assist in setting materiality that reflects your business and industry

News and press release sensing
Earlier awareness of emerging external factors that may impact risk of material misstatements that are relevant to financial reporting

Standardized risk reports
Clearer, more consistent insights for management and those charged with governance into potential risk facing the business and impacts on the audit

Draft walkthrough questions
A more efficient kickoff as we arrive with tailored questions to minimize disruption

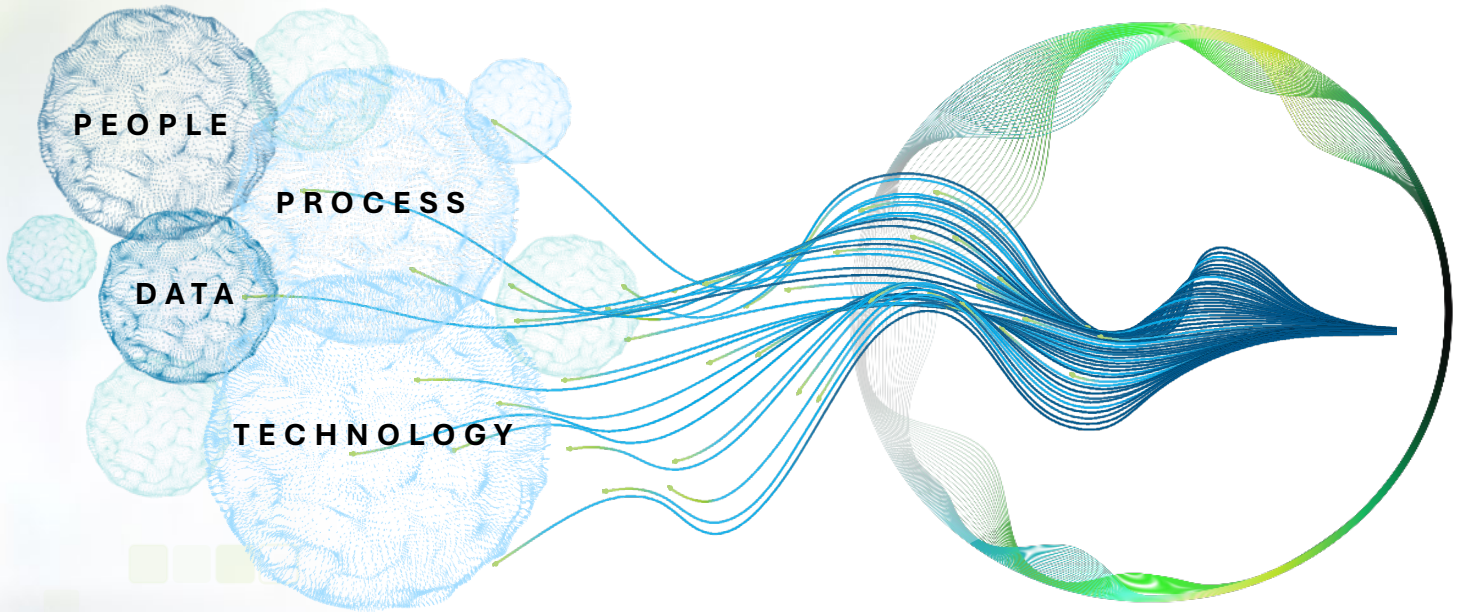
Key term extraction
Less time spent explaining documents or searching for details we identify what matters quickly

Financial statement analysis
Faster pinpointing of unusual movements or items that warrant early discussion

Copyright © 2026 Deloitte Development LLC. All rights reserved.

AI Risks to Consider

Extent of AI risk depends on the AI autonomy, complexity, extent of business integration, and maturity of organizational adoption.



AI RISKS (NOVEL/FAMILIAR)

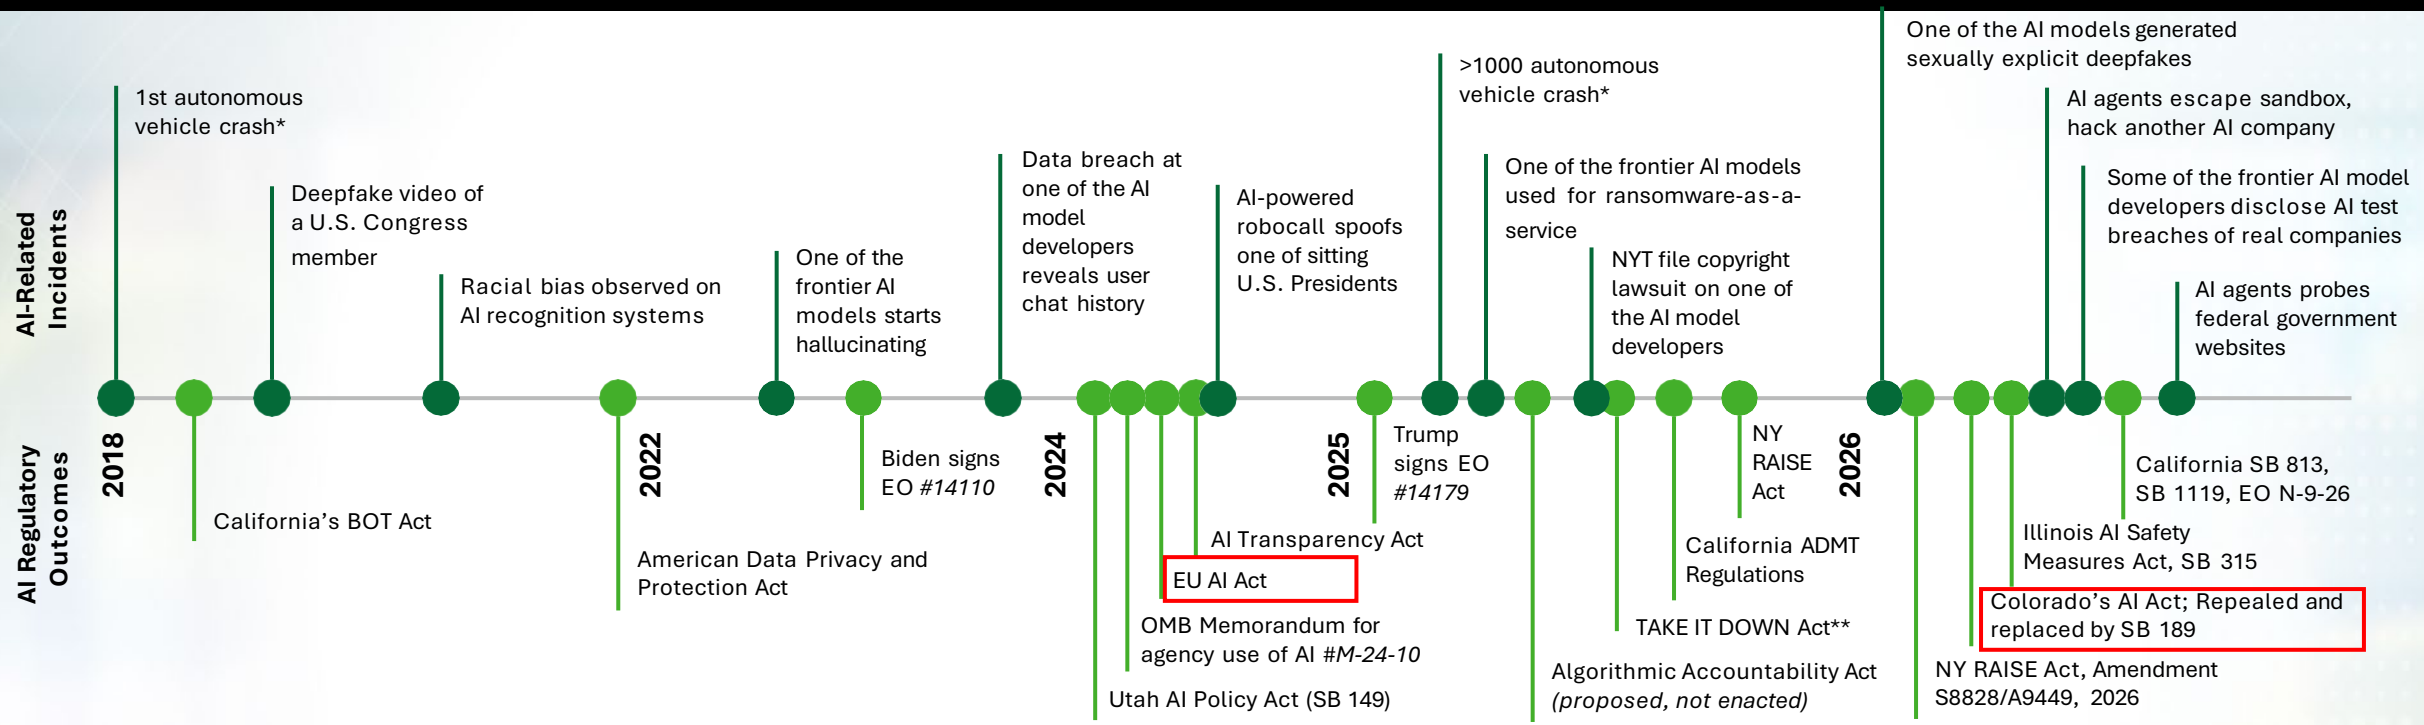
Explainability	IP infringement
Ethical use	Malicious intent
Bias	Confidentiality and privacy
Hallucination (Accuracy)	Access & Change Management

POTENTIAL CONSEQUENCES

Reputational damage • Destruction of shareholder value
Regulatory fines • Lawsuits • Poor strategic decisions

Copyright © 2026 Deloitte Development LLC. All rights reserved.

AI regulation is accelerating in response to real-world incidents…



Recurring Risks Driving Regulatory Action

Mass adoption of AI
78% of organizations report using AI in 2024 - up from 55% in 2023

GenAI risks
Misinformation, hallucinations, bias, IP disputes in GenAI results

Real-world incidents
Autonomous vehicle crashes, biased hiring algorithms, AI-enabled fraud

Algorithmic bias & discrimination
Documented harms in hiring, lending, housing, and policing

Privacy intrusions
AI models relying on large-scale personal data

Safety & security risks
Usage of AI medical devices and autonomous vehicles rising dramatically

Information integrity
Deepfakes, political manipulation, identity theft, cybersecurity vulnerabilities

National competitiveness
Without standards, U.S. risks falling behind the EU and China in shaping global AI norms

**Reported autonomous vehicle incidents - https://www.nhtsa.gov/laws-regulations/standing-general-order-crash-reporting*
*** Tools to Address Known Exploitation by Immobilizing Technological Deepfakes on Websites and Networks Act (TAKE IT DOWN Act)*

Copyright © 2026 Deloitte Development LLC. All rights reserved.

...but the state of regulation is highly fragmented across the U.S.

AI regulation in the U.S. is shaped by interacting federal and state approaches, influencing governance and operational decision-making

How U.S. AI Regulation Is Structured Today

Absence of a Singular Federal Law

There is no single federal AI statute; oversight relies on Executive Orders, OMB guidance, and agency enforcement using existing laws

Highly Fragmented State Regulations

States and local governments are rapidly enacting AI regulations, testing varied approaches across use cases, and collectively creating an increasingly complex and fragmented regulatory landscape

Continuing Federal Preemption

Federal–state preemption debates persist. Harmonization could reduce conflicts, but premature preemption risks regulatory gaps

Emerging Regulatory Themes

Despite fragmentation, common requirements are converging on transparency, risk management, fairness, human oversight/redress, privacy, and security

Implications for Organizations

Multi-Regulator Compliance Required

Expect multi-regulator scrutiny; build evidence-ready governance (inventories, impact assessments, documentation, monitoring) to satisfy diverse federal expectations

Rising Compliance Complexity

Patchwork compliance raises cost and execution risk; adopt scalable “highest-common-denominator” controls for hiring, credit, health, and other consequential uses

Regulation Uncertainty

Preemption uncertainty means designs must be adaptable; modular policies and contracts help pivot if federal baselines override or complement state rules

Preparedness for Future Baselines

Converging themes signal tomorrow’s baseline; embedding transparency, bias testing, oversight, privacy, and security now reduces retrofit costs and incident exposure

Copyright © 2026 Deloitte Development LLC. All rights reserved.

AI Governance and Oversight

Management, boards, and audit committees must work together to achieve effective AI governance and oversight

AI GOVERNANCE AND OVERSIGHT

BOARD OF DIRECTORS

Has **ultimate oversight responsibility** for the business and the company's strategy

AUDIT/OTHER COMMITTEE

- **Understands how AI is used** in the organization, including back-office and business operations, and the impact on financial reporting and other disclosures/measures
- **Oversees management's assessment and response** to AI-related risks including related internal controls
- **Reviews relevant metrics and trends** on the performance of AI systems within the organization as well as wider environment to understand how AI risks are evolving

MANAGEMENT AND INTERNAL AUDIT

Develops, implements, and evaluates use
(First line)

- Maintains the AI use-case inventory
- Designs, deploys, and operate systems
- Performs data assessment, testing, and documentation
- Monitors performance, drift, incidents, and control thresholds

Supports, challenges, and monitors
(Second line)

- Establishes AI policies, standards, and risk classification
- Reviews use cases and challenges risk assessments
- Monitors compliance and control effectiveness
- Aggregates risk reporting and escalates exceptions

Internal audit
(Third line)

- Provides independent assurance and advice
- Assesses governance, control design, and effectiveness
- Evaluates the completeness and reliability of AI inventory and reporting
- Audits lifecycle governance and verifies remediation

Copyright © 2026 Deloitte Development LLC. All rights reserved.

Leading Practices Throughout the AI Lifecycle

DESIGN PROCESS

- **AI Ethics training** programs for AI developers
- **Consult a diverse group of stakeholders** and gather a variety of use-case scenarios
- Conceptualize and **document** the system's intended use, objectives, and risks
- Hold sessions with stakeholders to **identify potential sources of adverse** outcomes and risks
- Understand latest **regulatory and legal requirements**
- Identify potential sources of risks (e.g., bias) and put into place **plans for mitigation**

DATA

- **Reputable, well-documented** data sources
- Use **representative data sets for** AI training purposes
- Use **least number of data features** that meet performance goals
- Assess data elements/features (including post-transformations) for **potential bias**
- Identify and monitor **data quality metrics**

DEVELOPMENT AND IMPLEMENTATION

- A **clear statement** of purpose to develop model in line with its intended use
- Drive model selection process with **objective criteria** and a structured suitability assessment
- Use a variety of **development tests** to evaluate the model type assumptions, stability, bias, impact of potential limitations, and behavior over a range of input values (sensitivity and boundary testing)
- **Back-testing** (in-sample and out-of-sample) and cross validation results
- **Strong documentation** that discusses approach, assumptions, limitations, risks, and related mitigation controls

TESTING AND REVIEW

- A **second line of defense** for review and AI oversight
- AI **oversight with authority to challenge** AI design and development
- **Review conceptual design** of model including data elements
- **Review developmental tests** and perform additional tests where needed

ONGOING MONITORING

- Regularly **evaluate model performance** using relevant metrics looking for changes or degradation. Define escalation/incident process.
- **Compare performance to relevant baselines** or expected performance indicators
- **Use visualizations** to identify patterns/trends easily
- Monitor **distribution of input data**
- Evaluate the model performance on subsets of data (e.g. **edge cases and demographic groups)**
- **Monitor** model retirement dates for 3rd party models and define change management triggers

USE

- Use a **human-centric approach** for decision-making to identify and report bias
- **End-user training** that includes considerations for ethics and potential harm

Governance/Oversight

- Roles and responsibilities
- Policies and procedures
- Value/ethics statement
- Public disclosures

Copyright © 2026 Deloitte Development LLC. All rights reserved.

Regulatory & Standards Landscape

SEC

Disclosure of material AI risks; accuracy of AI-generated metrics in filings

AIUC

Certification standard built for AI agents. Formal certification is issued through AIUC and accredited auditors

COSO ICFR

Internal control framework: control environment, risk assessment, monitoring

NIST AI RMF

Voluntary risk management framework: Govern, Map, Measure, Manage

EU AI Act

Risk-tiered obligations on AI systems; high-risk classification for some finance use

ISO 42001

AI management system standard: emerging certification path for AI governance

Monitor evolving AI-in-financial-reporting guidance from standard-setters and professional bodies.

Copyright © 2026 Deloitte Development LLC. All rights reserved.