INNOVATION, CYBERSECURITY, AND TECHNOLOGY (H) COMMITTEE

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

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.

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

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.

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

Please check all that apply:	
State Connected	☐
Operational	☐
Regulator Request	☒
Urgent Request *(Requires Immediate EPMO review)*	☐

Name of the Project/Initiative:	Cybersecurity Event Notification Portal Project
Project Sponsor(s):	Miguel Romero
Person Completing Proposal:	Koty Henry
Submission Date:	03/13/2025

1. **OPPORTUNITY STATEMENT: When answering this question, please describe:**
   - *Why do we need to implement this project?*
   - *What problem are we trying to solve for?*
   - *How do we know this IS a problem to solve for?*
   - *Have we ever done something like this before?*

   The current cybersecurity event notification process required by the Insurance Data Security Model Law #668 (MDL #668) is fragmented and inconsistent across the jurisdictions where it has been passed, adopted, and implemented. This is burdensome for licensees to navigate during a cybersecurity event, adding significant costs to an already expensive process, and increases legal risk due to inconsistent compliance expectations. Additionally, this complexity is a friction point for the industry as legislatures consider the adoption of MDL #668, as licensees must bear additional compliance requirements. Lastly, the Cybersecurity (H) Working Group has charges that would be imperiled by inaction or failure, particularly charges #3 and #8. Without the additional efficiency a central portal would create, developing guidance for cyber events and coordinating any resulting work (Charge #3) would remain overly complicated and slow. Finally, assisting with the passage and implementation of MDL #668 (Charge #8) cannot be reasonably achieved so long as the marginal cost of compliance with a jurisdiction's notification requirement remains high, a problem the central portal will address. In summary, as cybersecurity risks intensify, this lack of a unified, efficient reporting system hinders timely response, complicates compliance, and weakens stakeholder confidence.

   The Cybersecurity (H) Working Group has been working to align the various areas of MDL #668 and the cybersecurity event notification portal is the final piece. Already, guidance documents have been adopted to reduce the compliance burden of industry by aligning the efforts of departments of insurance as they enforce sections 6 and 7: the Cybersecurity Event Response Plan (CERP), and the Insurance Data Security Model Law Compliance and Enforcement Guide. While both have been

adopted and address necessary areas of convergence, the ability to centralize the reporting of cybersecurity events requires technology, not a guide. This project is that technology.

The basic technology underlying the portal, being able to store and receive highly sensitive information while limiting access to only the appropriate users and aligning with well recognized control frameworks and will include clear governance over security, requirements well within the NAIC's capabilities and expertise. This initiative aligns with the CERP adopted in March 2024 and supports the implementation of MDL #668.

2. **PROPOSED SOLUTION: When answering this question, please describe:**
   - *What is the proposed or "ideal" solution?*
   - *What are the intended outcomes?*
   - *Are there any "Hard" Dates to consider?*
   - *Are there any known risks to account for?*

The Cybersecurity (H) Working Group adopted the CERP on March 17, 2024, which guides state insurance regulators on responding to cybersecurity event notifications, required under Section 6 of MDL #668, from licensees. The CERP, however, can only unify the response to a notification by departments, not the underlying requirements faced by licensees to provide notification in the first place. The proposed solution is to develop a secure portal hosted by the NAIC to receive, manage, and track cybersecurity event notifications from licensed entities in states that have adopted MDL #668.

The portal would be built with a minimum of features, aiming for a high degree of uniformity within a licensee-directed notification system. To accomplish this, a set of notification questions that represent the requirements of all MDL #668 states will be developed into a single, standardized notification form (attachment 1). Data access will be highly limited to only those departments with an adopted version of MDL #668, and the responsibility for selecting those departments will be upon the licensee. Additionally, security concerns should be assuaged by an annual disclosure of a System and Organization Controls Reporting (SOC) 3 report by the NAIC. The SOC 3 report is the public version of the NAIC's SOC 2 Type II that is conducted annually, and their report looks at the Security Trust Services Criteria.

Lastly, as we develop experience additional opportunities may present themselves wherein the portal may be improved. Any such future endeavor will be done in consultation with our stakeholders and interested parties.

Key Features

- Licensee fills out a single, standard notification form, which may be updated as additional facts become known. A draft version of the standard form can be found in attachment 1, where additional details about uploading files can be found.
- Licensee directs notifications by selecting the departments to notify, which may be updated. Secure, role-based access for regulators, providing access only for those who have passed an equivalent to Section 6 of MDL #668.

- Regulators to have the ability to satisfy recordkeeping requirements by downloading completed records.

3. **KEY RESOURCES: When answering this question, please describe:**
   - *What resources/teams are required to deliver the solution?*
   - *What is the projected resource level of estimates for each affected area?*

   The success of a cybersecurity event notification portal relies on the collaboration of internal NAIC teams and regulators, in particular the Cybersecurity (H) Working Group, as well as collaboration with licensees, as warranted.

   Resource needs for this initiative have been identified as follows:
   - *NAIC technical and cybersecurity teams for development and hosting*
   - *Regulatory and legal input to align with state law variations*
   - *Business analysts and project managers to oversee implementation*
   - *Training and support staff for successful rollout and adoption*
   - *Infrastructure Investment (servers/cloud services, licences)*

4. **EXISTING ALTERNATIVES: When answering this question, please describe:**
   - *Do we have an in-house solution in place that COULD meet the existing need?*
   - *Is a new purchase/build absolutely required to meet the project needs?*

   While no single solution currently provides this fully tailored experience, the NAIC has tools and platforms that can enable an efficient and scalable build of such a solution.

5. **VALUE PROPOSITION: When answering this question, please describe:**
   - *Why should the organization invest in this initiative?*
   - *What value will be created by doing this?*
   - *What happens if we DON'T do this?*
   - *If applicable, please indicate the letter committee that adopted this project.*

   The portal is the final piece of a series of initiatives intended to harmonize and align the various cybersecurity requirements found in MDL #668. Once complete, the NAIC membership should be able to more easily pass MDL #668 nationwide, as legislatures will face fewer hurdles from industry as their problems with the law's expansion will have been solved. Further, the streamlined compliance and improved incident response that licensees will achieve will improve operational efficiency by reducing compliance costs and complexity. The portal will allow licensees the ability to submit information about cybersecurity incidents once within the security of the NAIC's systems rather than to multiple states, multiple times, through multiple platforms containing varying levels of security.

   Many issues will arise if this project is not pursued. Without a central solution there will continue to be duplicative regulatory efforts to investigate cybersecurity breaches and varying levels of security for submission of cybersecurity breach notifications across departments. Most importantly, failing to align requirements will make it difficult for MDL #668 to be passed in

additional jurisdictions, imperiling work on the Cybersecurity (H) Working Group's Charges #3 and #8.

6. **KEY SUCCESS METRICS: When answering this question, please describe:**
   - *What key deliverables will make this project a success?*
   - *How will we know when we are successful?*
   - *What key metrics will be used to define "Doneness"?*
   - *How will we test/validate metrics?*

   The key deliverables necessary to make this project a success would include developing a portal that:
   1. Meets both industry and regulator confidence requirements regarding data security and confidentiality.
   2. Allows licensees to submit cybersecurity event notifications through the portal to those jurisdictions that have adopted a version of Section 6 of MDL #668 that licensees determine should be notified.
   3. Allows each state regulator to receive all information necessary to meet the statutory reporting requirements of their jurisdiction via a standard form.
   4. Allows both regulators and industry participants to satisfy their recordkeeping requirements through secure downloads and robust logging to ensure auditability.
   5. Allows licensees to update the information within the notification, via the portal, as their investigation progresses.
   6. Notifies licensee-selected regulators via email when new information is available for viewing.
   7. Allows licensees to update which regulators who have adopted MDL #668 have access to the cybersecurity event notification, in keeping with their legal responsibilities.

   To measure the success of the proposed portal for cybersecurity event notifications, the following key metrics may be considered:
   1. Number of Cybersecurity Event Notifications: Track the total number of notifications received through the portal to assess adoption and usage.
   2. User Adoption Rate: Monitor the number of licensees and state insurance regulators using the portal and the growth rate of new users. *Maybe consider quantifying the metric by saying something like "**Adoption by at least 5 states in Year 1; goal to scale to 10+ within 2 years**"*
   3. User Satisfaction and Feedback: Collect and analyze feedback from users through incremental surveys to measure satisfaction, system uptime, and identify areas for improvement.

7. **CUSTOMER SEGMENTS: When answering this question, please describe:**
   - *What is the impact to the organization, members, another project or NIPR if this project is not approved?*
   - *Who is the customer/member? Who is your audience?*
   - *Who are the business owners/stakeholders that will be impacted by this?*

- *Who are the end users?*

Without a central solution there will continue to be duplicative regulatory efforts to investigate cybersecurity breaches and varying levels of security for the submission of cybersecurity breach notifications across departments. Most importantly, failing to align requirements will make it difficult for MDL #668 to be passed in additional jurisdictions, imperiling work on the Cybersecurity (H) Working Group's Charges #3 and #8.

The key stakeholders for this project are members and licensees. Departments of insurance whose legislature has passed a version of Section 6 of MDL #668 will have access to a central repository where every cybersecurity event notification they are legally entitled to is available. This will reduce work for departments as they review and track individual notifications and reduce duplicative and redundant responses. Further, for those states whose legislature has not passed a version of MDL #668 but would like to, the portal will reduce the regulatory burden experienced by licensees as more states adopt.

Licensees will accrue the immediate benefit of a repository as the burden for notifying multiple state departments of insurance will be dramatically reduced through the implementation of a standard form and a push-based, licensee-directed system. Currently, licensees must navigate multiple regulatory schemes for the notification of cybersecurity events, all with their own unique approaches to implementation. This fragmentation has led to a slow and ultimately expensive process for licensees.

8. **COST STRUCTURE: When answering this question, please describe:**
   - *What known costs are associated with the project (i.e., Software, Hardware costs, etc.)*
   - *Are there any 3rd party vendor costs to consider?*

Staffing Options & Estimate:

Below are Assumptions and Staffing Options/Estimates for the Cybersecurity Event Notification Portal Project EPMO proposal dated 2/11/26. Note that the timing of this effort could affect the quoted amounts and timelines based on staff availability.

Assumptions:

- Deliverable will be a single solution for all jurisdictions.
- Prototype using Appian and Design review with Working Group is completed prior to project approval by EPMO.
- Team is comprised of Product Owner, 3 Software Engineers and 1 Software Quality Engineer
- Prototype is for discussion purposes; the team will meet UI/UX and Appian guidance for the development work.
- Option one assumes the team is working 75% of the time on this project with the rest of the time allocated to other work to support existing applications.

The following phases are recommended:

Phase	Key Activities
Project Preparation Phase	Training (Option 2), planning, understanding project, refinement of work.
Development (MVP)	Build core features: intake form, role-based access (admin, company and regulator, audit trail, notifications). Reuse components and tables from UCAA and SERFF where appropriate.
Testing & Validation	Functional testing, automation testing, Dynatrace Synthetic, security validation, UAT with early adopters.
Pilot Rollout (5–10 states)	Controlled launch, feedback loop, refinements.
Full Rollout & Support Setup	Training, documentation, help desk, onboarding additional states.

- H Committee staff support will provide ongoing administrative support.
- On-going staff considerations also include .25 ITG FTE for maintenance and support post-release.

Staffing Options/Estimates

1. **A dedicated ITG delivery team with existing Appian development experience and no outside help.** This team will not need startup and training time/costs; however, this option may not be feasible based on anticipated workloads.

Cost: $0 consulting costs – NAIC internal labor only
Duration: 7.5-8 months

2. **An ITG delivery team with no Appian development experience and the assistance of one full-time NAIC Appian Software Engineer.** This approach enables the delivery team to complete current tasks within an extended duration. Training and startup time allocated is included for a team new to Appian. The mentor's team will have reduced capacity for the duration of the project.

Cost: $16,500 instructor led training and certification for 3 engineers, the team will be unavailable for other work.
Duration: 9.5-10 months

3. **Outsourced to a professional services consulting group, with oversight of NAIC staff.**

Cost: $2.1M consulting costs, and reduced capacity of NAIC staff working with outside firm
Duration: 7.5-8 months

9. **RISK MITIGATION: What risks should you mitigate to make the project successful? When answering this question, please describe known risk factors associated with implementing this project, such as:**
    - *Are there risks to achieving a high adoption?*
    - *Is this a new effort we've never done before?*
    - *Do we need to acquire innovative technology to implement?*

The NAIC is exceptionally well positioned to construct and administer the portal, and it fits well into its existing expertise. Accepting data, even highly sensitive and confidential data, is something NAIC does well and has done for a long time across many domains. However, some general risks remain, which are discussed below.

Risk: Stakeholders lack a method by which to understand the current security posture of the portal.

Mitigation: NAIC should make available, annually, a SOC 3 report for public review by licensees and other interested parties. The SOC 3 is the public version of a SOC 2, an industry-recognized and respected third-party assurance report, which preserves the sensitive security information of the NAIC.

Risk: Access and confidentiality protections for licensees.

Mitigation: Access to the portal will be based on whether a department's legislature has passed a version of Section 6 of MDL #668. The licensee, utilizing a push-based system, will select the departments to receive their notification as their legal responsibility requires. Like the form, the departments able to view the notification can be updated by the licensee. The NAIC will have limited access to enable support only, and will mirror the access they retain for other, highly sensitive information like RBC data, ORSA reports, and exam information.

Risk: Difficulty in tracking changes and updates to a notification by departments.

Mitigation: Given the simplicity of the design of the portal, with licensee-directed notifications, this difficulty is best addressed through a robust logging capability. This will provide the necessary insight for departments to understand how the notification has been changed and updated over time.

Risk: The underlying data increases in its sensitivity, which may result in higher cybersecurity risk to the licensee cybersecurity event information.

Mitigation: If highly sensitive information begins to be included in the portal, then cybersecurity risk becomes elevated in importance. To handle the prospective risk of elevated cybersecurity, a SOC3 report should be provided annually, which will ensure stakeholder clarity that highly sensitive information is adequately secured by NAIC.

Risk: Using the portal is difficult and causes problems for departments.

Mitigation: Clear instructions will be made available and training, if necessary, will be developed and offered.

While other risks may be present, all the risks found insofar have been found to have mitigation strategies that are acceptable to stakeholders.

10. **Member Support:** When answering this question, please document the members who are in support of this initiative.

Ensuring the alignment of the majority of members is crucial for the success of any initiative. The new EPMO project process is structured to foster collaboration and build consensus through a transparent, phased approach:

- The proposal will first be presented to the Cybersecurity (H) Working Group, providing a focused forum for subject matter experts and key stakeholders to evaluate the initiative, raise concerns, and offer recommendations. This early engagement ensures technical and operational considerations are addressed before moving forward.
- Following the Working Group's input, the proposal will advance to the Innovation, Cybersecurity, and Technology (H) Committee for broader member review. This step allows for strategic alignment across the organization to ensure that all members have an opportunity to provide feedback and influence the direction of the project.

This transparent, consensus-focused approach will provide the necessary alignment among all stakeholders to allow for member support to build organically.

11. **PROJECT DEPENDENCIES: Is this project dependent on other projects or initiatives? If yes, please list.**

No, this project does not depend on other projects or initiatives.

12. **HARD DEADLINES: Is there a deadline driving this project? ☐YES ☒NO If yes, what is it?**

13. **REVENUE STREAMS/SOURCES: When answering this question, please describe:**

- **Will this effort generate additional revenue or cost money to implement? ☐YES ☒NO**

Revenue generation from non-licensees may be explored in later phases. The current proposal is not predicated upon the recovery of costs from licensees.

- **If so, what is the revenue projection?**

14. Could an additional fee be charged to recoup costs and/or are there future budgetary cost savings? ☐YES ☒NO

15. Will NIPR share costs? ☐YES ☒ NO If yes, indicate your rationale and list the NIPR contact.

16. Provide high-level estimates for the initial and future costs associated with this project.

- **Please insert additional rows if needed.**

Software Licenses and/or Subscriptions – Type *(include maintenance on separate line)*	Number Requested	Individual Cost	Starting Month and Year	Length of Initial Term	% Allocated to NIPR
				3-Year	
		$			
TOTAL		$			

Hardware Purchases – Type *(include maintenance on separate line)*	Number Requested	Individual Cost	Starting Month and Year	Length of Initial Term *(maintenance only)*	% Allocated to NIPR
		$			
		$			
TOTAL		$			

Consulting – Staff Aug. Position Title (each requested consultant on its own line)	Proposed Hourly Rate or Fixed Engagement	Estimated Hours for Contract Term	Length of Contract Term	Starting Month and Year	% Allocated to NIPR
	$				
	$				
TOTAL	$				

Consulting – Prof. Services Position Title (each requested consultant on its own line)	Proposed Hourly Rate or Fixed Engagement	Estimated Hours for Contract Term	Length of Contract Term	Starting Month and Year	% Allocated to NIPR
TOTAL					

Training / Conferences / Certifications - Type	Number Requested	Individual Cost	Month and Year Attending or Start of Subscription	Is this an annual subscription. (Yes / No)	% Allocated to NIPR
		$			
		$			
TOTAL		$			

Travel – Purpose and Location if Known	Number Individuals Traveling	Number of Nights per Individual	Individual Cost from Current Travel Matrix	Month and Year of Travel
			$	
			$	
TOTAL			$	

New Headcount Requests – Job Description Title	Number Requested	Proposed Salary	Starting Month and Year
		$	
		$	
TOTAL		$	

17. What assumptions have been factored into the project estimates?

- It is assumed that the number of workforce identities will remain consistent at 1,000 over the 3-year term. While this number is currently accurate, the NAIC may experience growth.
- The estimates assume that the implementation and integration of the Identity and Access Management (IAM) tool with existing systems (Active Directory, Okta, eDirectory, Workday) will be completed within the planned timeline without significant delays.
- It is assumed that internal teams, including the IAM team and identity teams, will be available and have the necessary expertise to support the implementation and integration efforts.
- The estimates assume that the selected IAM tool will be compatible with the NAIC's existing identity systems and will not require significant additional customization or development.
- The cost estimates do not include a managed service offering for ongoing support. The NAIC may choose to purchase this offering in the future.

18. Please indicate the staff resources needed for this project in the table below.

Please insert additional rows if needed. Only technical hours will be tracked for the project.

Replace 0 with numbers. Highlight Total Number, Right click, Update field.

Internal Resources	Area/Team	Number/Type (Ex: 2-Analysts, 3-SE)	Total Estimated Hours

19. What is your confidence level in the above estimates? ***Low estimates will not be considered by the EPMO*****.**

☐ **HIGH**
Please comment:

☐ **MEDIUM**
Please comment:

20. Does the project include any of the following: PII/MNPI/other confidential information, attachments or ad-hoc data access? ☐**YES** ☐**NO**
If yes, please provide details regarding the confidential information, size, and number of attachments/retention period or ad-hoc data access needs.

-
-
-
-

For EPMO use only – do not fill out.

Account Description	Account Code / Dept	Total Expense for Initial Budget Year	Estimated Expenses for Following Year	Total Capital for Item (if Applicable)	Starting Month of Amortization or Depreciation	Length of Term	% Allocated to NIPR
		$	$	$			
		$	$	$			
Totals for EPMO spreadsheet		**$**	**$**	**$**			
Totals for NAIC		**$**	**$**	**$**			
Totals for NIPR		**$**	**$**	**$**			

Draft: 8/23/26

Big Data and Artificial Intelligence (H) Working Group
Columbus, Ohio
August 13, 2026

The Big Data and Artificial Intelligence (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met in Columbus, OH, Aug. 13, 2026. The following Working Group members participated: Nathan Houdek, Chair, and Coral Manning (WI); Doug Ommen, Co-Vice Chair (IA); Mary Block, Co-Vice Chair (VT); Richard Fiore (AL); Alex Romero (AK); Lori Munn (AZ); Chris Erwin (AR); Ken Allen (CA); Jason Lapham (CO); Joshua Hershman (CT); Nicole Crockett (FL); Matt Kilgallen (GA); Weston Trexler (ID); Ryan Gillespie (IL); Jake Vermeulen (IN); Saima Sanjida Shila (LA); Sandra Darby (ME); Mary Kwei (MD); Caleb Huntington (MA); Kate Stojsih (MI); Phil Vigliaturo (MN); Brad Gerling (MO); Connie Van Slyke (NE); Gennady Stolyarov II (NV); Sean Rosene (NH); Dave Wolf (NJ); Vanessa DeJesus (NM); Angela Hatchell (NC); Colton Schulz (ND); Matt Walsh (OH); Nicole Nash (OK); Gary Jones and Michael McKinney (PA); Andreea Savu (SC); Tony Dorschner (SD); Michael Schulz (TN); Marianne Baker (TX); Eric Lowe (VA); Ellen Potter (WV); and Lela D. Ladd (WY).

1. Adopted its July 22 Minutes

The Working Group met July 22 and took the following action: 1) received an update on the Artificial Intelligence (AI) Risk Evaluation Supplement pilot process and 2) heard a panel discussion on perspectives from actuarial organizations on generalized linear models (GLMs) and AI governance.

Stolyarov made a motion, seconded by Trexler, to adopt the Working Group's July 22 minutes (Attachment Two-A). The motion passed unanimously.

2. Received an Update on the AI Risk Evaluation Supplement Pilot

Commissioner Houdek said the 12 states participating in the pilot have been engaged since March and have met nearly weekly to discuss their individual state experiences implementing the pilot. He said that, based on feedback received during the pilot and input from subject matter experts, the supplement will be revised with an updated version ready for a public exposure comment period in early September.

He said each participating state has been engaged in the work over the spring and summer, and that some states used the supplement as part of a market conduct or financial examination that was already planned, while others used it as an ad hoc questionnaire outside the typical examination process. Each state has the flexibility to modify or add questions to meet its specific needs. The pilot will continue through September, and feedback has been collected from companies involved in the pilot, with each participating state also reaching out to its pilot companies to solicit input through a centralized survey of the experience. Written feedback has been received, which will help inform revisions to the supplement. Opportunities for verbal feedback will be provided during the public meetings following the two exposure periods planned for this fall, with the goal of finalizing the supplement and adopting the final version at the Fall National Meeting.

Crockett commented that the pilot process has been effective due to the continued engagement and strong collaboration between the states.

3. Heard a Presentation from AM Best on AI Governance in Insurance

Edin Imsirovic (AM Best) presented on AI governance in insurance, stating that the presentation was for discussion purposes and does not change AM Best's criteria, methodology, or rating guidance, nor does it create a formal AI governance assessment framework. He said the discussion on AI governance draws on public regulatory and supervisory standards, model risk, and AI security measures to address how governance expectations may change as AI systems move from predictive to generative to action-oriented workflows. He said the NAIC *Model Bulletin on the Use of Artificial Intelligence Systems by Insurers* and the AI Risk Evaluation Supplement pilot provide a foundation for a company's AI governance policy toward how governance operates.

Imsirovic defined three working terms: 1) predictive AI, which uses data to score, rank, classify, or flag something, often used in pricing models or claims triage, usually supporting a decision inside a well-defined process; 2) generative AI, which creates new content, such as summarizing documents, drafting language, extracting details, or turning unstructured data into something more useful, but its output can sound confident when it is wrong; and 3) agentic AI, which goes a step further and can work through a series of steps toward a goal, use tools, look up information, update a system, route work, or move a workflow forward. This system may affect processes directly rather than support a human decision. Each creates a different governance problem.

Imsirovic said AI does not create a separate risk category, as the exposures already flow through existing categories, such as product and underwriting, operational, and regulatory and legal risk, and that the enterprise risk management (ERM) question remains whether risk management capability is commensurate with the risk profile.

He said that well-governed AI deployment may support outcomes such as expense and processing improvement, underwriting quality, claims efficiency and consistency, and fraud control, while risks of execution, control breakdowns, cyber and data exposure, vendor dependency, and regulatory and explainability exposure can arise from deployment. He said what matters is whether governance mitigates the risks.

He said that if AI affects pricing, risk selection, segmentation, or portfolio mix, that points to product and underwriting risk. Data quality, claims-handling errors, cyber exposure, workflow breakdowns, or weak oversight of third-party models fall under operational risk and consumer protection. Explainability, litigation, or regulatory scrutiny can flow through the legislative, regulatory, judicial, and economic categories. One use case can span several categories; for example, a claims AI tool that begins as an operational issue becomes a regulatory issue if consumers are treated inconsistently and creates financial consequences if errors affect payments or reserve information. He said the categories are familiar, but the path into them is changing rapidly, as carriers move from predictive models to generative models to agentic actions. As a result, risks can emerge faster, become harder to understand, and move further without human intervention.

He described the development of AI capabilities as a compounding progression rather than a timeline. He said that using AI for fraud detection and pricing, or for computer vision and language processing used in claims and document work, was foundational, and that the governance problem at this stage remained closer to traditional data and model governance. He said that generative AI changed that, as those systems work with less structured information, produce less predictable output, and may be harder to explain, presenting a different governance problem rather than a harder version of an old one. He said that agentic AI changed it again, as those systems can act rather than only recommend, and one error can trigger others and move through a workflow before a person can intervene; therefore, a single governance approach is not enough.

He noted that predictive AI raises familiar questions about whether the model was properly tested, is fair, is documented, and whether its performance is changing over time. Generative AI raises questions about whether the output can be trusted, whether sensitive information stays where it should, and what happens when the underlying model is updated. Agentic AI raises questions about the system's access, what gets recorded, where a person steps in, and whether an action can be stopped or reversed. The risks build, such that an agentic system may still carry predictive and generative risks in addition to new permission and action risks, and that the comparison should be read as a comparison of governance problems rather than as a simple risk ranking.

He said that for predictive AI, monitoring should lead to action, so that when a model drifts or produces a weak outcome, the natural thing to look for is whether the company adjusted a threshold, limited the model's use, or added human review, and a dashboard that never changes anything is weak evidence. For generative AI, the question is whether controls hold up through the live workflow, including whether the level of review increases when the output matters more, whether the company tests the system again when a vendor updates the model, and whether the company can show that sensitive information is secured. For agentic AI, the evidence has to be more concrete, including whether permissions are tested, whether actions are logged for company reconstruction, whether there is a kill switch to roll the process back to a safe state, and whether a person has the authority to override actions. He said two companies can have the same policy but very different governance in practice.

He said the predictive, generative, and agentic labels do not indicate how much is at stake. He identified five factors: 1) how important the decision is, including whether the output directly affects a consumer or a financial outcome; 2) how much the system can do on its own; 3) whether the company can explain why the system produced a particular result or took a particular action; 4) how quickly the system can change, whether gradually through drift or suddenly through a vendor update; and 5) how dependent the insurer is on a third party, including whether the company can independently test, challenge, or override the system. Underlying these factors is proportionality, in that governance around a system should become stronger as stakes, autonomy, lack of transparency, speed of change, or third-party dependence increase, and he said this is consistent with the NAIC *Model Bulletin on the Use of Artificial Intelligence Systems by Insurers* and with recent work from the European Insurance and Occupational Pensions Authority (EIOPA), which also takes a risk-based and proportionate approach to AI governance in insurance. Thus, predictive AI is not inherently lower risk, and agentic AI is not automatically unacceptable. The label indicates the type of system, while the use case indicates more about the level of risk and the governance it may need.

He said lower-impact tools that mainly assist a person, such as an underwriting research assistant, still need governance, but the likely consequences of an error are more limited, and a person remains close to the process. Where the system is still assisting a person but the output matters more, such as extracting information from medical records for underwriting or preparing claim summaries for adjusters, the review has to be meaningful. This raises the question of whether the reviewer has enough information, time, and authority to question the result rather than accept it. He said that where the system has more autonomy but the impact is operational, such as narrowly defined automated endorsement processing, the questions become whether the boundaries are clear, whether actions are logged, and whether errors can be contained and reversed. He said the most demanding governance zone combines greater autonomy with more important consumer or financial outcomes, such as an agentic system routing and paying claims or making underwriting decisions. Controls such as human override, action logging, rollback, and clear limits on the system's actions become especially important.

He said controls are built in layers, with earlier controls staying in place as each layer is added. Predictive AI governance begins with model risk controls, including independent validation. Revalidation triggers mean deciding in advance which kinds of changes should lead to another review or testing, such as changes to a predictive model

or its data. For generative AI, a company should know what prompts are used; otherwise, the output will be harder to explain. For agentic AI, the level of approval should depend on the action, as looking up information is not the same as changing coverage or issuing payment. Some controls depend on vendor contracts, such as testing rights and notice of important updates, which can be much harder to obtain.

He said that where governance lives can make a difference, as two companies can use the same AI tool and have similar controls but end up with very different levels of risk in how the tool is integrated into the business. Governance should be built into the process. Fragile tool layering occurs when a tool is added to an existing process without changing the workflow or its accountability. He said companies do not choose that setup deliberately but arrive at it gradually. As a pilot works, more people begin using it, the use case spreads, and only later does the company realize that the process around the tool never caught up. A well-designed predictive model in a weak operating environment can create more risk than a more advanced AI system that is properly integrated and controlled.

He said shared dependency is broader than reliance on a single vendor, as many carriers may depend on the same vendors, models, data sources, cloud providers, or other infrastructure, and noted that the Financial Stability Board (FSB) has highlighted third-party dependency and provider concentration as potential financial stability concerns. For a single carrier the issue comes down to three questions: 1) visibility, whether the company understands how the system behaves; 2) control, whether the company can test and challenge the system, be notified when something important changes, and step in when needed; and 3) portability, whether it is difficult to replace or unwind a tool once it becomes part of an important workflow. A vendor change can affect how a system behaves, even if the carrier makes no changes, and risks can compound if the company does not review it. The larger point is at the market level, as each carrier may see only its own dependency, but if many carriers rely on the same underlying provider, a single weakness can affect many companies at the same time. This does not mean third-party use is a problem, as vendors can offer useful tools, but the key is whether insurers can understand what the system is doing and have the authority to oversee it and step in when something changes. The issue is whether oversight works effectively.

He identified four areas where a company should be able to show more than a policy or framework: 1) AI inventory, where knowing which decisions each system touches and how much it can do on its own shows which systems need the most attention; 2) validation and monitoring, where the question is what happens when a trigger is hit, including any tests and approvals performed before the system went live and what the company did when monitoring later found a problem; 3) human oversight, where a company should be able to show an example in which a person challenged the result, changed it, or stopped the process; 4) vendor and agentic controls, where the question is whether the company can see a vendor change, test its effect, stop the system, or return to a safe version. Policy alone is not enough; rather, operational governance is necessary.

He said the question is whether the use of AI stays bounded, monitorable, governable, and resilient as systems change, take on more autonomy, and become harder for people to oversee. There is still no single AI rulebook, but regulators want to know what AI is being used for, what decisions it affects, how it was tested, who is accountable, and evidence showing that controls actually work. He said that the last point represents the real shift, as much of the earlier work was principles-based, while newer work is moving toward implementation and examination.

He said modern AI systems may include the model, the prompt, the information it retrieves, memory, tools, and permissions, so the control boundary becomes much larger than the model itself. He said the UK AI Security Institute (AISI), working with the Bank of England, examined what developers are building rather than only testing models in a laboratory, reviewing more than 177,000 agent tools built over the preceding 16 months. Two findings

stood out: 1) tools that take action by changing something in an external system grew from about one-quarter of usage to 65% of usage over that period; and 2) while most action tools support medium-stakes work, such as editing a file (finance was an exception), with high-stakes financial roles having more tools that can change things than would be expected, including growth in the number of servers that can execute payments from about 46 to more than 1,200 in a single year. He said memory fits within that same inventory question, raising data governance questions about what is kept when a system carries information from one task to the next, how long it is kept, and whether it can be corrected or deleted when a policyholder asks, with the difference being that the data now sits inside the system rather than beside it, so identity, delegated authority, the ability to revoke access, and what the system remembers may all become core governance questions rather than technical details.

He said that capability can change even when the model does not. For example, the AISI tested frontier models on multistep cyberattack scenarios and increased the compute budget from 10 million tokens to 100 million tokens, which improved performance by 59% with no retraining, additional fine-tuning, or special expertise. This means a company could keep the same model and still materially change the risk by changing what surrounds it, whether through more time, more attempts, better tools, or wider permissions, such that approving the model may not be the same as approving the deployment.

Commissioner Houdek asked where companies struggle the most in governance. Imsirovic said that one of the biggest areas for improvement is building or revising the operating model around the use of AI, as too many companies use various tools without redesigning the operating model.

Block commented that agentic AI seems like a transition from human to AI models, back to another human being doing the work. Because agentic AI operates independently, it feels like a continuum in which supervision of a nonhuman employee is required. She appreciated the discussion of having a plan for a quick kill switch, for what happens when a system fails, and for the process to shut a system down, and that she views this as part of the ERM disaster recovery process to be treated the same way as planning for when a key employee disappears. Imsirovic agreed and said that agentic AI is not as widespread as generative AI at the moment, in part because companies are cautious about deploying agentic systems due to operational and regulatory concerns. He said most of the deployments he sees are in lower-stakes roles, such as a frequently asked questions (FAQs) chatbot or research assistance, and that the controls companies most often discuss involve keeping a human in the loop throughout the process and ensuring the AI system is monitored and tested, with the caveat that agentic AI remains mostly in early pilot stages.

Commissioner Ommen said a number of Working Group members are receiving different responses through that structured effort. He asked whether Imsirovic is beginning to see policies and procedures standardized in writing or whether there is still a great deal of communication through meetings. Imsirovic responded that the amount of capability an AI system has draws a different amount of scrutiny within insurance companies, noting that companies typically tell him they do not use AI for underwriting or, if they do, that it is monitored, and that in underwriting, it is usually used to turn unstructured data into structured data to enrich the submission process. Another area is in the claims process, helping adjusters with claims summaries, and that the most consistent answer he receives to governance questions is that a human is in the loop. Commissioner Ommen asked whether that response comes in conversation or whether written procedures are beginning to appear. Imsirovic responded that most of it comes through conversations, as AM Best does not have an official process for gathering data on AI governance, although that may change. He said some companies volunteer more information than others and that most are keen to describe how AI will improve their processes, but he has not seen a great deal of that documentation.

Stolyarov asked whether Imsirovic is seeing insurers adopt a consequential approach to AI governance in which the impact of a particular failure is considered not only for a consumer but for the enterprise as a whole. He illustrated the distinction with a chatbot that interacts with consumers and hallucinates information, which could mislead consumers and draw regulatory concern but is fundamentally correctable because a person at the insurer could intercede and provide the correct information, compared with an AI agent that, in an effort to be helpful, proactively and irreversibly deletes the company's claims information to free storage space, which is a failure from which the enterprise as a whole could not recover. The member cited the July 2024 CrowdStrike incident, in which a new employee implemented a systemwide update that disabled the systems of airlines and other enterprise clients, as the kind of catastrophic failure governance needs to prevent, and asked whether insurers are prioritizing those systemwide consequences. Imsirovic responded that the level of sophistication depends in part on the sophistication of a company's use of AI systems, as companies that entered the space earlier had more time to develop proper governance frameworks for those tools and tend to be higher-rated with well-established ERM frameworks. He said his concern is with companies that are newer to the technology and lack that history, which may implement it without having thought it through.

Commissioner Hershman noted Imsirovic had identified three types of AI, that those systems are leveraged for different use cases carrying different risks to the consumer or the company, and that different governance structures should surround those different types of systems and risks. He asked how the Working Group should think about the different levels of governance and how governance tools should be distinguished, such as by high risk versus low risk or by the type of AI being used. Imsirovic responded he would look for a pattern rather than one document, including a useful AI inventory, clear testing and approval standards, monitoring that leads to action, examples in which human review changed the outcome, evidence that the company can respond when vendor systems change, and whether the company is changing its operating model around the technology. He said the answer returns to the impact chart and the five factors: The higher the impact of a particular tool, the more governance that tool will need. He said an agentic tool used for narrowly defined endorsement processing may not require as much governance as the same type of tool used to pay claims or to automate underwriting, so the use case and what is at stake should determine the level of governance around a particular technology.

4. Discussed Other Matters

Commissioner Houdek reminded participants that an updated version of the supplement will be exposed for a public comment period at the beginning of September.

Having no further business, the Big Data and Artificial Intelligence (H) Working Group adjourned.

SharePoint/NAIC Support Staff Hub/Member Meetings/H Cmte/2026_Summer/WG-BDAI/Minutes-BDAIWG081326-Final.docx

Draft: 8/6/26

Big Data and Artificial Intelligence (H) Working Group
Virtual Meeting
July 22, 2026

The Big Data and Artificial Intelligence (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met on July 22, 2026. The following Working Group members participated: Nathan Houdek, Chair, and Coral Manning (WI); Doug Ommen, Co-Vice Chair (IA); Mary Block, Co-Vice Chair (VT); Alex Romero (AK); Mark Fowler (AL); Tom Zuppan (AZ); Ken Allen (CA); Jason Lapham (CO); Wanchin Chou and George Bradner (CT); Nicole Crockett (FL); Weston Trexler (ID); Victoria Hastings (IN); Shaun Orme (KY); Caleb Huntington (MA); Raymond Guzman (MD); Sandra Darby (ME); Kate Stojsih (MI); Phil Vigliaturo (MN); Brad Gerling (MO); Connie Van Slyke (NE); Sean Rosene (NH); Randall Currier and Howard Wegener (NJ); Vanessa DeJesus (NM); Gennady Stolyarov (NV); Nishtha Ram (NY); Matt Walsh (OH); Nicole Nash (OK); Michael Humphreys (PA); Elizabeth Kelleher Dwyer (RI); Andreea Savu (SC); Haelly Pease (SD); Michael Schulz (TN); Rachel Cloyd (TX); Eric Lowe and Dan Bumpus (VA); and Lela D. Ladd (WY).

1. Adopted its June 1 Minutes

The Working Group met June 1 and took the following action: 1) received an update on the artificial intelligence (AI) systems evaluation tool pilot process and 2) heard a panel discussion on AI governance trends.

Fowler made a motion, seconded by Ladd, to adopt the Working Group's June 1 minutes (Attachment Two-A1). The motion passed unanimously.

2. Received an Update on the AI Systems Evaluation Tool Pilot Process

Manning provided an update on the AI systems evaluation tool pilot process. She said the pilot officially began in March and involves 12 participating states, whose insurance regulators continue to hold weekly calls to learn from each other, raise questions for the group, and receive training and presentations. She said that, in response to feedback from many stakeholders about the name of the tool, it has been renamed the AI Risk Evaluation Supplement to reduce confusion about what the supplement is and is not.

She reviewed the timeline for the remainder of the pilot through the Fall National Meeting. She said the individual state pilots have continued through June and July at varying speeds and that education and discussions have continued. The pilot will continue through September, and companies selected for the pilot should look for opportunities to provide input, including a centralized survey of the pilot experience, with each state coordinating its own invitations to companies individually. The next version of the supplement will be shared at the end of August, followed by two rounds of public exposure to provide ample opportunity for public input, in addition to feedback from participating companies, with the goal of adopting the supplement at the Fall National Meeting.

3. Heard a Panel Discussion on Perspectives from Actuarial Organizations on GLMs and AI Governance

Commissioner Houdek said the Working Group developed a draft version 4.0 of the supplement to help regulators better understand the scope and uses of AI systems and assess the potential risks of those systems as they are used in insurers' operations, including the data used to train them. He said one topic that interested parties and regulators have repeatedly raised is whether generalized linear models (GLMs) should be included or excluded from the scope of the supplement, as GLMs have been used in insurance company ratemaking for decades, and a

robust regulatory framework already exists for the review of GLMs used in ratemaking. However, GLMs are also used in other insurer operations, which have not been fully addressed by the current regulatory framework. He said the Working Group has also discussed the desire for an AI governance program to include components such as transparency and requirements for effective human-in-the-loop practices.

Commissioner Houdek asked the panelists about issues the Working Group should consider regarding AI governance and AI-related risks from an actuarial perspective.

R. Dale Hall (Society of Actuaries—SOA) said that insurers and actuaries are approaching AI through governance and risk management, digesting the NAIC *Model Bulletin on the Use of Artificial Intelligence Systems by Insurers*, along with existing corporate risk management frameworks. He noted that the SOA Research Institute participates in the U.S. Department of Commerce's AI consortium through the National Institute of Standards and Technology (NIST), whose AI risk management framework is widely discussed across the insurance industry, and said that it is useful that the supplement asks how AI systems integrate with a company's Own Risk and Solvency Assessment (ORSA) and enterprise risk management (ERM) assessments. He said actuaries view most AI tools as a new variation on models the profession has used for decades, so questions of data quality review, validation, and documentation are not entirely new. He cited Actuarial Standard of Practice (ASOP) No. 56 as a comprehensive guide for designing and evaluating models.

Henry Liu (Casualty Actuarial Society—CAS) said the risk that is most top of mind is bias, both in the statistical sense emphasized in actuarial training, which focuses on balancing models so that results regress toward the average, but also in a social sense, in that models should not introduce preferential treatment or injustice into business processes, which becomes easier to miss as the model becomes more complicated.

Spencer Sadkin (American Academy of Actuaries—AAA) said that bias relates to the actuarial fairness and explainability questions the profession has examined over the past 30 to 50 years and cautioned that overly prescriptive governance frameworks become unadaptable. He noted that AI and AI systems carry distinct definitions and urged frameworks built on strong principles rather than prescriptive requirements. Commissioner Houdek responded that the Working Group has discussed the same concern from a regulatory standpoint, given how quickly the technology is changing.

Commissioner Houdek asked the panelists to describe GLMs and how they differ fundamentally from other machine learning and generative AI algorithms.

Sadkin said that GLMs can be thought of as a more naturally explainable approach and that a fundamental difference is that whoever creates the GLM specifies the structure up front, before the model is created, unlike most other machine learning algorithms, where the algorithm is selected and then allowed to solve and produce the answer. He said generative AI differs in that it creates something as it goes, whereas the other models are supervised learning approaches intended to answer a specific question.

Liu said that GLMs are methods that assume linear, or curvilinear, relationships between the inputs and the output variable. GLMs have many real-world applications, but the constraints of this simple model form become apparent when relationships are more complex, which is where more advanced model types and training techniques come in. He said that a GLM is strictly a linear model form, whereas AI encompasses both model forms and modern computational techniques for training. He added that all of these models are attempts to describe relationships observed in the real world and that correlation does not mean causation.

Hall said that the starting point assumption of GLMs is that variables are fundamentally assumed to have a linear relationship, and that in an insurance context, GLMs allow actuaries to explicitly specify the input factors, called rating variables, that are assumed to have a mathematical relationship with an outcome for the determination of rates or premiums. He said linear regression dates back roughly 200 years and that GLMs were introduced on the actuarial risk modeling exam taken after the probability, statistics, and financial mathematics exams, where candidates learn to build, validate, and explain them. He said that rather than viewing GLMs and generative AI as entirely separate categories, there is a spectrum, with GLMs among the more transparent, easier to explain, and mathematically based approaches.

Stolyarov commented that GLMs do not fit the definition of AI because they are static, deterministic models in which humans completely prescribe the assumptions, the model structure, and the input data, so that the output is essentially determined by those human inputs, whereas true AI works in a way that cannot be precisely predicted by humans in advance. He referred to the definitions in the NAIC private passenger auto survey on artificial intelligence and machine learning circulated in 2021, which describes what AI and machine learning systems include and exclude, and said the distinction needs to be made more prominently as generative AI becomes more common. He illustrated the point, noting that a spreadsheet or a GLM performs computations in an exactly correct manner arithmetically, whereas a large language model may return a calculation accurate only to a few significant digits unless prompted to perform a precise computation, in which case it might call code capable of producing a precise result. He continued to state that emerging generative models require a different set of approaches and precautions to have confidence that the models are performing basic math correctly.

Liu responded that, within a universe of AI models, some are more unpredictable than others, but that the variability of large language model output is partly a commercialization choice by the builder, such that the model form is not inherently random but can be more flexible. He added that a GLM trained on a data set as large as those used for large language models would require computational and algorithmic simplifications, in that training on a sampled subset would not guarantee the same outcome, so a taxonomy based on the observation that GLMs feel more deterministic would ignore that GLMs can also be made more unpredictable depending on the data used for training. Sadkin commented that the question of whether GLMs are considered AI is the wrong discussion and does not matter, as academia considers GLMs to be machine learning techniques, and the distinction makes no difference to the end consumer affected by the models. He said the more important discussion concerns the different levels of scrutiny that may need to be applied based on the end user's familiarity with a model and how it is used, and he noted that the Academy is working on a paper on this topic.

Eric Ellsworth (Checkbook Health/Center for the Study of Services—CSS) asked whether there are clear definitions of when non-deterministic outputs are acceptable and when they are not. He said whether a model is AI may be a semantic question compared with what guardrails are required, expressed concern about the non-determinism of large language models, and asked how the process of understanding or mapping variability differs for a large language model compared with other approaches and at what point the logic of a process should be frozen and clear, compared to LLMs which can be fluctuational. Sadkin said the approaches used will differ depending on the use case, and that an actuary building a pricing model is almost always using a supervised learning technique in which the modeler selects the algorithm, often a GLM because GLMs naturally resemble rating plans, after which the model is static and has been assessed up front. He said the generative AI use cases he has seen in insurance are more often tasks such as compiling the 10 to 20 notes a claims adjuster may have into an overall summary, where the human-in-the-loop concept applies because the adjuster can review the individual underlying entries, reducing the need for stronger governance processes.

Liu added that the objective underlying the discussion of non-deterministic answers is avoiding catastrophic failures, such as a claim denial not based on solid evidence, and that human-in-the-loop means having someone

with human judgment and experience to prevent those failures, but any human process is also prone to randomness, as different claim examiners adjusting an identical claim may reach different outcomes, so it is important to be realistic about the level of human randomness already present in the business process that AI is replacing.

Hall said the matter comes down to asking good risk-evaluation questions, such as what data was used for training, how it was validated, what documentation exists, and what transparency can be offered. He said validating the use of a GLM is almost like turning in math homework, in that the data, the calculations, and the predictive value can all be shown, whereas layering broader generative AI topics on top calls for showing how the model was trained and that its output is auditable, transparent, reliable, and consistent across repeated runs, so that the appropriate request is to show the risk evaluation performed rather than only the mechanics of the mathematics.

Commissioner Houdek said transparency is frequently discussed by regulators, industry, and consumer representatives, but the definitions of these terms, including AI and AI systems, differ across these conversations. He asked what transparency should consist of when incorporated into a governance program.

Hall responded that transparency includes the model's purpose, its data sources, its methodology, and the validation performed, including identifying contexts in which a company is and is not comfortable using the model, and how the model is monitored and its change history to incorporate new information and retesting. He said that in an industry with protected classes, those types of questions make sense to avoid bias.

Sadkin said the ultimate concern behind transparency is protecting the consumer and challenged the view that GLMs are highly transparent while other techniques are opaque. He said that examining a decision tree shows it is entirely transparent because at each point the path goes either left or right. Such models are highly transparent, and diagnostics on variable importance features can be available on the back end.

Liu agreed that GLMs have an overhyped reputation for transparency, saying that once a model has more than roughly a dozen variables and 100,000 training records, a coefficient can be observed, but the reason it takes a particular value cannot be explained. He said there are both model-specific and model-agnostic diagnostics, such as a lift chart, which does not depend on which model generated the output, and that many of the ways regulators already evaluate GLMs have model-agnostic equivalents that may be more suitable for wider adoption, so that any model can be governed and made transparent.

Lapham agreed with the distinction between transparency toward consumers and the transparency regulators expect, but the two can bleed together because, as consumers become more aware that insurer decisions are informed by these tools, they come to regulators seeking an explanation, and it will be up to regulators to help them sort through it. He said documenting the changes made to a model over time is extremely important to transparency for regulators, and that it is also important to understand not only that a third-party modeling firm or the insurer has performed good model governance, but also, as has been observed in Colorado, when a carrier has decided not to use a model or not to use it for a specific purpose, and the rationale for that decision.

Liu responded that the transparency of the data feeding into a model is as critical as the choice of model, recalling that two decades ago, the use of credit in insurance was discussed, with reporting indicating that credit information could be inaccurate roughly 30% of the time. He said that feeding poor data, or data generated by vendors' AI systems, into a GLM is just as lacking in transparency and accountability as the concerns raised about AI within an insurance company, and he offered the example of an insurer using satellite imagery to assess the shape of his own roof, noting that he would want to know whether the assessment used his roof rather than his

neighbor's and would want the opportunity to validate and contest it. He said that as models take in more data, that data should also be subject to scrutiny and appropriate due diligence.

Commissioner Houdek relayed a question asking whether, given the potential for the characteristics of GLMs and AI to blend, there could be a model that is GLM-like but leaves it to AI to find and use at least one data endpoint in the service of a desired goal, such as minimizing risk in specific circumstances or keeping costs below a certain level, without further guardrails. Liu responded that such approaches exist because an AI model is not necessarily the most cost-effective option for a specific business process, so a simpler model may be used in production while AI is used to train on the back end. He said these are likely not being used for ratemaking, as he has not seen filings requesting approval of a process to derive rating values on the basis that whatever the process produces will be used, but that outside of ratemaking, AI-trained GLMs and other simple model forms are in use today within insurance and within the general economy.

Sadkin cautioned that discussions of AI should be crystal clear in distinguishing between AI models and algorithms, as the term appears opaque at present, and noted that the European Insurance and Occupational Pensions Authority (EIOPA) has discussed the same question of whether GLMs should be carved out and developed proposals on how that should work. He said his view, which he believes others share, is that such an approach expends unnecessary additional regulatory effort to categorize modeling techniques. Rather, it should be based on the level of potential risk, rather than on establishing a framework that allows matters to flow through naturally. Commissioner Houdek said the NAIC participates in the EU-U.S. Insurance Dialogue Project, an informal joint project with the EIOPA, in which AI regulation and supervision have been a focus area for the last couple of years.

Having no further business, the Big Data and Artificial Intelligence (H) Working Group adjourned.

SharePoint/NAIC Support Staff Hub/Member Meetings/H Cmte/2026_Summer/WG-BDAI/2026 0722Interim-Meeting/Minutes-BDAIWG072226-Final.docx

Draft: 6/10/26

Big Data and Artificial Intelligence (H) Working Group
Virtual Meeting
June 1, 2026

The Big Data and Artificial Intelligence (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met June 1, 2026. The following Working Group members participated: Nathan Houdek, Chair, Timothy Cornelius, and Coral Manning (WI); Doug Ommen, Co-Vice Chair (IA); Mary Block, Co-Vice Chair (VT); Alex Romero (AK); Mark Fowler (AL); Lori Drever Munn (AZ); Ken Allen (CA); Jason Lapham (CO); Wanchin Chou and George Bradner (CT); Karima M. Woods (DC); Ron Waye (FL); Shannon Hohl (ID); Ryan Gillespie and Ricardo Mancilla (IL); Shaun Orme (KY); Caleb Huntington (MA); Raymond Guzman (MD); Sandra Darby (ME); Kate Stojsih (MI); Phil Vigliaturo (MN); Brad Gerling (MO); Colton Schulz (ND); Connie Van Slyke (NE); Christian Citarella (NH); Randall Currier (NJ); Gennady Stolyarov II (NV); Nishtha Ram (NY); Matt Walsh (OH); Nicole Nash (OK); David Dahl (OR); Michael Humphreys and Joe Schwartz (PA); Elizabeth Kelleher Dwyer (RI); Andreea Savu (SC); Travis Jordan (SD); Michael Schulz (TN); Nicole Elliott (TX); Eric Lowe (VA); Curt Kwak (WA); and Lela D. Ladd (WY).

1. Adopted its Spring National Meeting Minutes

Stolyarov noted two corrections to the Spring National Meeting minutes.

Colton Schulz made a motion, seconded by Stolyarov, to adopt the Working Group's March 24 minutes (*see NAIC Proceedings – Spring 2026, Innovation, Cybersecurity, and Technology (H) Committee, Attachment One*) with the two corrections proposed by Stolyarov. The motion passed unanimously.

2. Received an Update on the AI Systems Evaluation Tool Pilot Process

Commissioner Houdek noted that the Working Group would provide interim updates between the Spring National Meeting and Summer National Meeting to keep stakeholders apprised of progress.

Manning stated that the pilot officially started in March 2026 and involves 12 participating states. All states have sent, or are in the process of sending, inquiries to their respective companies. The group of pilot state regulators continues to hold weekly calls to share insights on anticipated responses and has received multiple training sessions on data science, compliance, and governance concepts related to the tool. Training will continue throughout the summer. Pilot state regulators have largely identified the companies that will receive inquiries and, in many cases, have already notified those companies. When a company is part of a larger group, pilot states coordinate to avoid duplicate requests.

Manning said that participating states are moving at varying paces. Some have already sent Exhibit A and/or Exhibit B and are reviewing responses, and at least one state has progressed to Exhibit C. Other states have not yet sent inquiries, as they are synchronizing the exhibits with examination timing. Implementation methods include formal examinations, separate surveys or data calls, or a combination of both. A survey has been developed and will soon be deployed to all participating carriers, with responses open through Sept. 30. Two more draft version of the tool will be exposed for public comment prior to the Fall National Meeting.

Manning noted that key areas identified for revision include tightening definitions, including clarifying the distinction between "AI program" and "governance program" in Exhibit A to ensure consistency with the *Model Bulletin on the Use of Artificial Intelligence Systems by Insurers*; tightening materiality language; and clarifying

Exhibit A information requests. The tool has been effective in allowing states to tailor their use to individual companies. She outlined the timeline for the remainder of the pilot process through the Fall National Meeting. The pilot will continue through September, with official pilot-state feedback closing after that. Company surveys will be open through Sept. 30. Tool version 5.0 will be drafted using feedback received through Aug. 1 and introduced at an end-of-August Working Group meeting, then exposed for a 30-day public comment period throughout September. In October, the Working Group will review comments on version 5.0, pull company survey responses through September 30, and draft version 6.0, which will be exposed for 14 days. In early November, quick revisions will be made, and a Working Group meeting will be held the week before the Fall National Meeting to address comments received on version 6.0. Version 7.0 will be presented at the Fall National Meeting for adoption consideration. An additional Working Group meeting with actuarial panelists is scheduled for July 22.

Lindsey Stephani (National Association of Mutual Insurance Companies—NAMIC) commented that industry is receiving similar feedback that definitions could use additional clarity, the scope of the tool should be tightened, and that it is encouraging that materiality is being considered. She asked whether the Working Group is examining the definition of AI and whether the potential inclusion of predictive or rules-based models is within scope. Manning responded that the pilot group is currently in the collection phase, gathering comments from individual states before conducting broader internal discussions. She indicated that specific definitional feedback submitted in writing would be considered in the next draft.

3. Heard a Panel Discussion on AI Governance Trends

Commissioner Houdek welcomed panelists Michael David Ruiz (Deloitte), Fred Karlinsky (Greenberg Traurig), and Mary Jane Wilson-Bilik (Eversheds Sutherland) and moderated a discussion on AI governance trends.

In response to a question about AI governance practices that are working well, Wilson-Bilik stated that companies are making meaningful progress in building governance and risk management programs that incorporate AI. She identified the following effective practices: establishing board- and senior-management-level AI accountability structures; developing strong data governance policies to ensure data is accurate, representative, and free from bias; monitoring outcomes for bias and hallucinations; and preparing AI inventories. She is seeing significant success in developing cross-functional governance teams and integrating AI into compliance and vendor governance programs. She noted that addressing human-in-the-loop requirements, particularly in the context of agentic AI, remains a difficult and emerging challenge.

Ruiz echoed the importance of cross-functional governance teams, noting that effective governance brings together stakeholders from three lines of defense, where operations, IT, technology, and data are brought to the table at the same time as the second line of defense of risk and compliance and legal, and then internal audit as the third line of defense. He stated that a good starting point is to establish a clear AI vision and philosophy from the top of an organization and translate them into policies aligned with the NAIC AI principles of transparency, explainability, security, and responsible use.

Karlinsky stated that lessons from the cyber governance space, including board-level ownership, cross-functional engagement, and treating AI as a transformational business issue, are being successfully applied in the AI context. These are transformational issues for companies, and they need to ensure they are addressed at the highest level and involve a broad cross-section of people working in the space.

He noted that the NAIC AI model bulletin, adopted by 25 states, provides a strong foundation for governance. He identified risk-based guidance, transparency, and explainability requirements, human-in-the-loop requirements, and proactive bias testing with annual attestations (as implemented in Colorado and New York) as emerging best

practices. He stated that the most effective governance programs create guardrails without gridlock, enabling innovation while protecting consumers.

Commissioner Houdek asked about areas where the language in the model bulletin or the AI Systems Evaluation Tool may be unclear or difficult to apply. Karlinsky stated that the NAIC's definition of AI has drawn comment from some in the industry, with some preferring the National Institute of Standards and Technology (NIST) definition. He also pointed out some confusion about the tool's context and use in a market conduct examination. He noted that while the tool is a solid starting point, technology regularly outpaces regulation and legislation, and that continued iteration, developed through regulator and industry collaboration, is necessary.

Wilson-Bilik stated that she has heard from companies that some questions included in the tool for financial examinations appear more oriented toward market conduct or consumer protection, raising the question of whether a single tool can serve both purposes. She identified materiality and the sometimes interchangeable use of the terms "model" and "system" as areas requiring definitional clarity. She also stated that the growing use of agentic AI raises significant new regulatory concerns, including AI agents exceeding their authority, cascading errors in downstream systems, real-time data governance challenges, and new cybersecurity risks. She suggested that both the model bulletin and the tool may need to specifically address agentic AI, noting that the bulletin provides a solid foundation but does not yet account for the systemic risks posed by agents that can operate beyond the boundaries of a single underwriting or marketing tool.

Ruiz agreed, stating that a comprehensive AI inventory should distinguish among generative AI, agentic AI, and simpler rules-based tools, and that a risk-based tiering approach applied consistently across the full AI lifecycle is the appropriate framework.

Karlinsky agreed that this would be a good first step, but companies may not be able to identify all the AI models. However, the industry is collectively gaining experience over time. The important part is that the industry is working more closely with the regulatory community to benefit consumers.

Commissioner Houdek asked what factors make a model or AI system high risk. Wilson-Bilik identified the following: the intended use of the model; the sensitivity and accuracy of the data on which it is trained and which it analyzes; the model's ability to interpret context (particularly in chatbot applications); the model's potential adverse consumer impacts; and potential downstream and cascading systemic effects. She mentioned the recent Mythos model developed by Anthropic as an example of the power and danger of applications of AI.

Karlinsky cited vendor risk, model drift, proxy discrimination, black-box models, data governance, and velocity mismatch as key risk-focus considerations, noting that many are analogous to issues previously encountered in credit scoring and cyber regulation.

Ruiz stated that insurance companies are implementing risk-tiering models across the full AI lifecycle, from use-case identification through deployment and ongoing monitoring, treating the initial risk rating as the "building permit" for a model's development and continued operation.

Eric Ellsworth (NAIC Consumer Representative) raised the challenge from a consumer perspective of fragmented accountability across business units and vendors, i.e., the "nobody home" issue, where consumers cannot easily reach a person, such that accountability is fragmented across business units and vendors, with no clear person having enough knowledge of what the system is doing to solve their problem. He asked the panel to discuss maintaining organizational knowledge of process steps and gears, evaluation criteria, and availability of an empowered human to solve problems, as models and agents are adopted.

Wilson-Bilik referenced the Colorado AI Act (SB 26-189), which requires companies to designate a trained individual with the authority to override consequential automated decisions, and noted that this type of "meaningful human review" standard (defined as someone who is trained in the workings of the of the actual automatic decision making technology [ADMT]) could serve as a model for addressing consumer accountability.

Karlinsky acknowledged that the ultimate goal is for models to perform well enough that continuous human intervention is not required, while cautioning that it is premature to assume that point has been reached.

Ruiz stated that tracking the right key performance indicators as model capabilities evolve and reporting those metrics to governance stakeholders is essential.

Stolyarov asked whether insurance organizations have implemented sandbox or test environments for AI agents that can operate but cannot permanently or irreversibly alter any system, and then, when the AI agent provides its recommendation, it is escalated to a human reviewer for action. The consequences of the AI agent in this environment would not flow through to the organization, the customers, or the general public unless a human being says this is an acceptable result.

Ruiz replied that he has not yet seen a seamlessly deployed test-to-production human-in-the-loop interaction in practice, but noted that the concept is actively developing and that the proliferation of agentic AI will demand its further formalization within governance frameworks.

Karlinsky commented that he has seen a lot of testing to get the right answers using AI, but believes it is still in its early stages and that there will be more testing. Wilson-Bilik stated that rigorous testing prior to production deployment is essential and represents sound governance practice, but that it is still too early to be able to know how this will play out.

Miguel Romero (NAIC) asked the panelists to confirm his understanding that the way a model is used and the data underpinning it are the foremost drivers of model risk. Wilson-Bilik agreed, adding that potential consumer outcomes and financial impacts must also be incorporated into the risk assessment, including the model's training. Ruiz agreed, adding that the criticality of the use case and the breadth of affected stakeholders, particularly where model outcomes directly impact consumers in areas such as underwriting and claims decisions, are equally important dimensions of model risk. Karlinsky agreed that consumer impact is key, but model drift and how models are applied are equally important.

Ellsworth commented that the most important thing for consumers is being able to accomplish the tasks they are interacting with the company to accomplish, and often that involves crossing many system boundaries. He asked whether the panelists had heard of any emphasis on testing types for accountability across systems, because the natural breakdown is for someone to own each one and someone to think about each one. Karlinsky replied that a cross-disciplinary approach is something he has been seeing for a while. Ruiz added that comprehensive model inventories should document system dependencies, interconnectedness, and cross-stream and downstream impacts, with designated model owners and stakeholders responsible for those dependencies.

4. Discussed Other Matters

Commissioner Houdek noted that overall engagement from pilot state regulators, companies, and industry throughout the pilot process has been very positive. Romero confirmed that the next Working Group meeting is scheduled for Wednesday, July 22, 2026, at 11:00 a.m. CT and is posted on the NAIC website. The July 22 meeting will include an update on the AI systems evaluation tool pilot process and a panel discussion with actuarial panelists.

Having no further business, the Big Data and Artificial Intelligence (H) Working Group adjourned.

SharePoint/NAIC Support Staff Hub/Member Meetings/H Cmte/2026_Summer/WG-BDAI/2026 0601Interim-Meeting/Minutes-BDAIWG060126-Final.docx

Draft: 7/7/26

Cybersecurity (H) Working Group
Virtual Meeting
June 10, 2026

The Cybersecurity (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met June 10, 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 (AR); Kurt Swan (CT); Tim Li (DE); Paula Shamburger (GA); Daniel Mathis (IA); Shane Mead (KS); Nina Hunter (LA); Mary Kwei (MD); Danielle Torres (MI); Bubba Aquirre (MN); Kim Dobbs (MO); Joshua Hilliard (NH); Gille Ann Rabbin (NY); Matt Walsh (OH); David Buono (PA); Curt Kwak (WA); and Barbara Belling (WI).

1. Heard a Presentation on the NetDiligence 2025 Cyber Claims Study

Peterson introduce Mark Greisiger (NetDiligence), who is the president of NetDiligence, and provided background on his experience in cyber risk assessment, cyber insurance, and data breach services. Peterson noted that presentation materials had been distributed to the Working Group members in advance. Greisiger presented findings from the NetDiligence 2025 Cyber Claims Study, drawn from approximately 4,000 new claims contributed by roughly 150 partner cyber insurers, with a five-year look-back dataset of approximately 10,400 claims. He explained that the study focuses primarily on small and mid-sized enterprises (SMEs), with an average SME revenue of approximately $108 million and an average incident cost of approximately $203,000 in 2024, compared with an average incident cost of approximately $5 million for large entities (averaging approximately $12 billion in revenue).

Greisiger described ransomware as a leading peril and a top driver of cyber insurance demand. He reported that initial ransom demands against SMEs averaged $3.1 million and were typically negotiated down by approximately two-thirds with the assistance of carrier-engaged forensic experts and breach counsel, resulting in average payments of approximately $1 million. He shared observations from NetDiligence's Ransomware Advisory Board, a quarterly gathering of approximately 70 leading forensic and cybersecurity firms, which included: 1) increasing threat-actor bypass of SMS-based multifactor authentication (MFA), with authenticator-app-based MFA identified as a stronger control; 2) circumvention of endpoint detection and response tools through misconfiguration and social engineering used to disable controls; and 3) the prevalence of help-desk impersonation attacks. He reported that advisory-board polling on extortion case volume shifted from approximately 54% reporting "stable" volume in the first quarter (with approximately 30% reporting increases) to approximately 70% reporting increasing volume in the second quarter.

Greisiger reviewed sector-specific findings, noting that professional services (including accounting and law firms) led incident frequency with an average cost of approximately $520,000, followed by manufacturing at approximately $811,000, and technology and healthcare each averaging approximately $1.4 million per incident. He identified business email compromise (BEC) as the second-leading cause of loss, typically manifesting as Automated Clearing House (ACH) or wire fraud through executive impersonation, and reported average SME business-interruption losses of approximately $630,000.

Greisiger explained that nearly all 2024 claims reviewed in the study were associated with criminal activity, representing a significant shift from prior years when system glitches and employee errors accounted for a larger percentage of losses. He further noted that “nano” SMEs (revenue at or below $50 million) represented approximately 41% of breached entities in the dataset. He observed that losses of unencrypted laptops due to theft or loss have become increasingly uncommon due to widespread encryption practices. He also stated that

ransomware actors increasingly exfiltrate all available business data, making traditional distinctions between protected health information (PHI), personally identifiable information (PII), and other categories of information less meaningful in claims analysis.

Greisiger highlighted emerging artificial intelligence (AI)-driven exposures, including threat-actor use of AI-based voice and video filters that enhance social engineering and impersonation across calls and videoconferences. He anticipated future AI-related cyber liability claims (including copyright and trademark infringement, bias, and defamation arising from AI-generated content) and observed that AI liability will likely also touch directors and officers, errors and omissions, products, and employment practices lines. He noted that NetDiligence engages with more than 200 cybersecurity vendors that are also leveraging AI technologies to improve defensive cybersecurity capabilities.

Greisiger offered to share with the Working Group: a full copy of the current claims study; the forthcoming 2026 study upon its release; a sanitized version of the carrier data-collection form; and to explore potential Working Group access to NetDiligence's policyholder portal which supports peer benchmarking by sector and entity size.

Peterson asked what data elements would be necessary to advance from monitoring claim volume and base-level categorization toward understanding incident impact. Greisiger described approximately 10 coverage "buckets" reflected in the carrier collection framework, spanning first-party categories (such as business interruption and revenue loss) and liability categories (including breach counsel, forensic response, and regulatory enforcement defense, such as state attorney general actions). He added that NetDiligence increasingly supplements carrier-supplied data with input from forensic partners, which often provides more granular details (such as starting demand versus negotiated payment, and root-cause safeguard failures) than claims departments can provide.

Peterson asked how the adequacy of cyber insurance coverage relative to actual incident costs is currently assessed. Greisiger responded that organizations frequently purchase limits that are insufficient to cover the full costs associated with significant cyber incidents. He noted that insurers can offer substantial coverage limits and that brokers may construct layered excess coverage towers when additional capacity is needed. However, he explained that forensic, legal, and incident response expenses incurred immediately following an event can consume policy limits before investigations, recovery efforts, and remediation activities are complete.

In response to a question regarding the basis of claim sampling, Greisiger explained that NetDiligence's process is relationship-driven, with partner carriers contributing redacted claim data through a standardized collection form and noted that the methodology is described in the full report's methodology section.

Grassel asked how NetDiligence distinguishes criminal from non-criminal events. Greisiger explained that ransomware and BEC events involving wire or ACH fraud are categorized as criminal by definition, while other events are classified based on the underlying claim description. He noted that future categorization will be complicated by the emergence of AI-related cyber liability.

Norman Niami (American Property Casualty Insurance Association – APCIA) asked about the methodology used to develop the claims dataset and how trends are validated over time. Greisiger explained that participating insurers voluntarily provide anonymized claims data and that additional methodological details are available in the full study report. He noted that NetDiligence relies on long-standing relationships with participating insurers and cybersecurity service providers to support data collection efforts.

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

NAIC Support Staff Hub/Member Meetings/H Cmte/2026_Summer/WG-Cybersecurity/2026 0610Interim-Meeting/Minutes-Cyber061026.docx

Draft: 8/18/26

Privacy Protections (H) Working Group
Columbus, Ohio
August 13, 2026

The Privacy Protections (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met in Columbus, OH, Aug. 13, 2026. The following Working Group members participated: Elizabeth Kelleher Dwyer, Chair (RI); Alex Romero (AK); Mark Fowler (AL); Damon Diederich (CA); Kristin Fabian (CT); Jordan Esbrook (IA); Victoria Hastings (IN); Mary Kwei (MD); Sandra Darby (ME); Colton Schulz (ND); Connie Van Slyke (NE); Roger Hayashi (NV); Tyler Trammell (OK); John Haworth (OR); Frank Marnell (SD); Dan Bumpus (VA); and Bryan Stevens (WY).

1. Discussed Comments on Revised *Privacy of Consumer Financial and Health Information Regulation* (#672)

Director Dwyer stated that on July 24, the Working Group exposed the full draft of the revised *Privacy of Consumer Financial and Health Information Regulation* (#672) for a 60-day public comment period ending Sept. 22. She then asked for initial comments on the draft.

Bumpus (VA) expressed support for the revisions to Model #672 and stated that the draft reflects a compromise. He cited improvements to consumer notices, consumer rights related to access, correction and deletion of nonpublic personal information, limited opt-in rights, and preserved exceptions for service providers and joint marketing.

Rep. Steve Elkins (MN) said he had followed the initiative as a state legislator and as a member of privacy-related legislative working groups. He stated that the draft appeared to build on the Washington privacy framework reflected in comprehensive consumer data privacy laws adopted by multiple states, and he said he appreciated the Working Group's harmonization with existing state privacy laws.

Brent Walker (InsuranceFraud.org), discussed antifraud considerations. He said that while the draft contains some antifraud exemptions, he would identify opportunities to include further exemptions throughout the model in his written comments.

Wes Bissett (Independent Insurance Agents & Brokers of America—IIABA) stated that most state comprehensive privacy laws only apply to large businesses. He asked for a limited exemption to Section 5—Third-Party Service Provider Arrangements, for small businesses. Bissett also suggested the Working Group would be better served by directly regulating third-party providers.

Director Dwyer asked Bissett how the Working Group could place direct requirements on third-party service providers, given the lack of regulatory authority over them. Bissett responded that the legislature could provide enforcement authority to attorneys general or other regulators, including insurance departments.

Kirsten Wolfford (American Council of Life Insurers—ACLI) said the ACLI supported modernization and consumer protections while requesting careful review of definitions, preservation of the Gramm-Leach-Bliley Act (GLBA) opt-out framework, workable state models, and further opportunities to comment.

Lindsey Stephani (National Association of Mutual Insurance Companies—NAMIC) stated that NAMIC has concerns related to opt-in portions of the model, consistency with other state privacy laws, and deviations from the GLBA.

Ken Klein (Individual Consumer Advocate) argued that opt-in consent provides stronger privacy protection than opt-out. He recommended a general presumption requiring opt-in consent when information is shared for purposes beyond the consumer-requested insurance transaction.

Harold M. Ting (Individual Consumer Advocate) requested prohibitions on dark patterns, clearer notices and consent forms, more transparent categories of information, and stronger explanations of consumer choices.

Kristin Abbott (American Property Casualty Insurance Association—APCIA) emphasized foundational definitions; downstream effects; and sufficient implementation time for systems, notices, governance, vendor contracts, and compliance controls.

Director Dwyer requested detailed written comments and proposed language by Sept. 22. She stated the Working Group hoped to complete the project by December, subject to review of comments and scheduling.

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

SharePoint/NAIC Support Staff Hub/Member Meetings/H CMTE/2026_Summer/WG-Privacy/Summer-Minutes/Minutes-PPWG081326.docx

Draft: 8/19/26

Third-Party Data and Models (H) Working Group
Columbus, Ohio
August 12, 2026

The Third-Party Data and Models (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met in Columbus, OH, Aug. 12, 2026. The following Working Group members participated: Jason Lapham, Chair (CO); Nicole Crockett, Vice Chair, and Richie Frederick (FL); Chad Bennett and Alex Romero (AK); Richard Fiore (AL); Lori Dreaver Munn (AZ); Ken Allen and Chandara K. Phanachone (CA); George Bradner and Wanchin Chou (CT); Doug Ommen and Jordan Esbrook (IA); Weston Trexler (ID); Julie Holmes (KS); Matthew Stewart (LA); Jackie Horigan (MA); Raymond A. Guzman (MD); Sandra Darby (ME); Phil Vigliaturo (MN); Brad Gerling (MO); Colton Schulz (ND); Gennady Stolyarov (NV); Kevin Yan (NY); Stewart Trego and Matt Walsh (OH); Michael McKenney (PA); Beth Vollucci (RI); Andreea Savu and Will Davis (SC); Marianne Baker (TX); Jessica Baggarley (VA); Mary Block (VT); and Timothy Cornelius (WI).

1. Adopted its July 16 and Spring National Meeting Minutes

Lapham stated that the Working Group met July 16. During this meeting, it took the following action: 1) discussed the draft Third-Party Framework, PC Pricing, and Underwriting Data and Models, which it exposed July 8 for a 28-day public comment period ending Aug. 5.

The Working Group also met June 29 and April 9 in regulator-to-regulator session, pursuant to paragraph 3 (specific companies, entities, or individuals) of the NAIC Policy Statement on Open Meetings) to discuss third-party regulatory framework drafts.

Commissioner Ommen made a motion, seconded by Trexler, to adopt the Working Group's July 16 (Attachment Five-A) and March 22 (*see NAIC Proceedings – Spring 2026, Innovation, Cybersecurity, and Technology (H) Committee*) minutes. The motion passed unanimously.

2. Discussed the Third-Party Regulatory Framework for P/C Pricing and Underwriting

Lapham said the Working Group exposed the third-party regulatory framework for property/casualty (P/C) pricing and underwriting July 8 for a public comment period ending Aug. 5. Eleven comment letters were received.

Scott Harrison (American InsurTech Council--AITC) said AITC is not wholly supportive of a registry and would prefer alternative ways achieve the objectives of transparency, accountability, etc. He identified four concerns with the registry: 1) it would create a new bureaucratic infrastructure; 2) the benefit to the third-party vendor is unclear; 3) the registry would be public and identify the third parties with which insurance carriers are doing business; and 4) an existing review process for P/C rating and underwriting appears to be working and could be used by other states.

McKenney and Stolyarov both said it should not be problematic to have a public registry; the fact that insurers are doing business with particular third parties in rating and underwriting is public information in the states. The proprietary part of models is what is considered confidential. McKenney noted that with 7,000 P/C rate filings in Pennsylvania, the insurance companies are not privy to third parties' data sets and model details. He said the point of this framework is to create an avenue for regulators to go directly to third parties and get information.

He said it seems that alternative actions would require the adoption of a law by the legislature, where the outcome might not be uniform. Instead, state insurance regulators can implement this voluntary registry, work together, and get access to the information needed.

McKenney added that regulators want to have confidence that insurance carriers understand the models they are using and are using them appropriately in compliance with state laws and regulations. He said the registry would allow third parties to keep information confidential from insurers while still providing information regulators need. He said a level of transparency would be achieved.

Frederick said Florida has issues similar to Pennsylvania's. State insurance regulators would need to educate third parties to understand what they need to do to do business in a state. Transparency is needed for insurers too, at least to the extent that insurers would have a level of confidence and understanding about the information or tools they are using.

Kristen Abbott (American Property Casualty Insurance Association—APCIA) said this framework was a meaningful improvement over the original version, and APCIA appreciates that it has been narrowed to focus only on the P/C pricing and underwriting data and models. APCIA encourages a tailored, risk-based approach to any future expansion of the framework and further refinement of definitions and registration requirements to avoid unnecessary complexity and duplication. While the registry is described as voluntary, APCIA questions whether it would function that way in practice. It believes participation should remain truly voluntary and that insurers should not lose access to valuable third-party tools simply because a vendor chooses not to participate or disclose proprietary information to state insurance regulators. It also encourages continued focus on operational and implementation issues, including governance and reporting expectations, transition periods, treatment of legacy models, and situations where insurers depend on information that may only be available from vendors. Abbott said any final framework should recognize the practical limits of insurer oversight and support good faith compliance efforts. APCIA also appreciates the stronger confidentiality protections in the revised draft and encourages additional safeguards to proprietary vendor information and commercially sensitive insurer information.

Commissioner Ommen said this work has been ongoing for some time, and Iowa state insurance regulators have been investigating what can be done in Iowa. He said third-party vendors have been filing as advisory organizations. However, they are not really advisory organizations, so as a regulator, Commissioner Ommen said he is left with the question of whether to revoke or deny their registration as an advisory organization, to require them to withdraw from their registration, or to look for an alternative. He said APCIA appears to view the proposed voluntary process as a constructive alternative to state insurance regulators navigating varying state laws to identify a to receive filings. Abbott said feedback from members suggests this is a good starting point, with some further clarifications requested.

Tyler Hobbitzell (Blue Cross Blue Shield Association—BCBSA) said if the Working Group expands this work into other lines, BCBSA requests that it be evaluated independently with a stakeholder input. The final P/C product should not be presumed to apply to other lines. Lapham said that is the plan.

Laura Panesso (ISO Verisk) focused on clarifying the scope of regulated entities, avoiding requirements that may deter participation, and removing provisions that might duplicate other processes. She said "licensee" was defined in the first draft of the framework and would like it reinstated because it would distinguish regulated entities from third-party vendors and provide needed clarity regarding the definitions. The annual attestation's Section 3B addresses the insurers' compliance with applicable law, which is not controlled by third parties who would sign

the attestation. She said use cases for third-party data and models may vary. Different statutes may apply to those different use cases, and the vendor may not always know the exact use case that a particular insurance company customer is implementing. Panesso suggested softening the requirement so that it is reviewed in conjunction with applicable laws rather than in compliance to provide some flexibility for those various use cases.

Lindsey Klarkowski (National Association of Mutual Insurance Companies—NAMIC) said the Working Group should avoid regulatory duplication and confusion by revising its charges to coordinate directly with the committee working groups that are also addressing third parties. Most recently, the Market Conduct Regulation Modernization (D) Working Group has been having active discussions at the Summer National Meeting. While the efforts of this Working Group are very well intentioned and an important exercise, NAMIC is concerned that the proposed frameworks are duplicative of current insurer responsibilities and risk creating confusion given the overlap with existing filing rules and existing filing practices that various states have adopted. NAMIC believes that efforts across the NAIC should be coordinated, and this Working Group should look at existing rules and practices, identifying if there are gaps, what they are, what is working well, and what will help to avoid this result.

Commissioner Ommen said Iowa accepts filings put in the square box of "advisory organization." However, he said they are not advisory organizations, and that is the experience in numerous states. He asked whether NAMIC had an alternative comprehensive approach. Klarkowski said the frameworks were walked through, but NAMIC recommends digging into them to identify pain points. She mentioned Arizona recently implemented a new advisory organization-type category specifically for data and model vendors. States are allowing nonlicensed entities to also file in the System for Electronic Rates & Forms Filing (SERFF) as an informational filing. NAMIC wants to avoid the growing patchwork of approaches that are happening and respects the effort to try and have some cohesiveness there. Klarkowski said she thinks the conversation needs to start with an evaluation of what is and is not working in the states. Chou asked NAMIC to submit a specific list of duplications. He said the lack of an NAIC framework has led states to be creative, and state insurance regulators want to address that patchwork to create something more cohesive across states.

Block said this project began after industry and third-party vendors raised concerns with regulators about the lack of a uniform framework and the resulting challenges across all 50 states. She said the Working Group needs to not relitigate how it started this in the first place and solve this problem.

Klarkowski said NAMIC members have not reported issues getting information from the third parties. She said a communication pathway is needed between third-party vendors and state insurance departments because, under the current structure, states do no provide consistent confidentiality protections for third parties that submit proprietary information to departments, and the formal relationship exists only between the third party and the insurer.

Bryan Rehor (ZestyAI) said that, as a third-party property analytics provider, ZestyAI supports most of the draft and believes a single NAIC-hosted registry is an improvement over the current state-by-state patchwork. The opportunity for vendors to file documentation directly with regulators is the right design because it gives them complete access while protecting the underlying methodology. He said ZestyAI endorses governance expectations tied to the NAIC Principles on Artificial Intelligence (AI) and *Model Bulletin on the Use of Artificial Intelligence Systems by Insurers*, which are standards that responsible vendors should already meet.

Rehor raised three points from his letter. First, a vendor officer would not be able to attest, within an annual attestation, that the data and models comply with insurance law in every state where an insurer uses them. He said the vendor does not control how a carrier implements a score, which rating plan it goes into, or where it is

deployed. Second, confidentiality with open records policies is a concern. The draft encourages states to extend trade secret protection, but it acknowledges that some states will need new statutes. Vendor filing obligations should not attach in a state unless that state provides protection equivalent to that given to insurers. Third, Section 2b allows a state insurance regulator to bar all insurers statewide from using a vendor's models. There should be some materiality threshold, a written notice, and a cure period.

Chou said the industry and vendors have a significant number of staff that state insurance regulators do not have, and that is why regulators are asking vendors to take ownership and sign off on what they have developed via the attestation. Rehor agreed with the need for attestation but said vendors should also be required to do the same because vendors do not control the final implementation of the products. He said a model solution can be deployed in a perfectly compliant manner, and that same solution can be deployed in a noncompliant manner. Vendors can attest that the solution licensed to the subscriber is compliant. Chou agreed that insurers have their own ownership and responsibility.

Lapham said the plan is for the drafting group to make revisions for consideration by the Working Group. Ideally, the plan is to have a finalized framework ready for consideration at the Fall National Meeting, but he said the Working Group will need to see how the work progresses over the next three or four months.

Having no further business, the Third-Party Data and Models (H) Working Group adjourned.

SharePoint/NAIC Support Staff Hub/Member Meetings/H Cmte/2026_Summer/WG-Third-Party/Summer-Minutes/Minutes-Third-PartyWG081226-Final.docx

Draft: 7/30/26

Third-Party Data and Models (H) Working Group
Virtual Meeting
July 16, 2026

The Third-Party Data and Models (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met July 16, 2026. The following Working Group members participated: Jason Lapham, Chair (CO); Nicole Crockett, Vice Chair (FL); Charles Hale (AL); Kara Voss (CA); George Bradner and Wanchin Chou (CT); Doug Ommen (IA); Weston Trexler (ID); Julie Holmes (KS); Jackie Horigan (MA); Raymond A. Guzman (MD); Sandra Darby (ME); Phil Vigliaturo (MN); Brad Gerling (MO); Gennady Stolyarov (NV); Kevin Yan (NY); Matt Walsh (OH); David Dahl (OR); Beth Vollucci (RI); Nicole Elliott (TX); Dan Bumpus (VA); Mary Block (VT); and Nathan Houdek (WI).

1. Discussed the Third-Party Regulatory Framework for P/C Pricing and Underwriting

Lapham said the Working Group exposed the draft of the regulatory framework on July 8 for a public comment period ending Aug. 5, with comments to be compiled for continued discussion at the Summer National Meeting. He said the current draft represents phase one of a broader effort, as the Working Group originally considered a broader scope, including claims handling, utilization review, marketing, and fraud detection. However, after receiving comments and deliberating, the Working Group reached a consensus to narrow its initial focus to third-party data and models used in property/casualty pricing (P/C) and underwriting.

Lapham said this focus reflects the prevalence and consequence of third-party vendor models in these areas and states' existing experiences reviewing vendor filings in this context, and that other functional areas remain on the table for future phases.

Lapham said the executive summary sets out three purposes for the framework: 1) compliance, to help states confirm that insurers' use of third-party data or models complies with existing state insurance law and does not produce rates or underwriting results that are excessive, inadequate, or unfairly discriminatory; 2) governance, to encourage robust data and model governance standards on the part of vendors; and 3) visibility, to give state insurance regulators comprehensive awareness of, and direct access to, the third-party data and models used in ratemaking and underwriting decisions. He said that regardless of the level of oversight applied to third-party vendors through registration, attestation, or filing requirements, none of it shifts responsibility away from the insurer, which remains fully responsible for compliance with all applicable insurance laws when using third-party data or models, and vendor registration is not a substitute for an insurer's own due diligence.

Turning to definitions, Lapham said the framework distinguishes between models and data. A pricing model is a simplified representation of how losses, expenses, and risk factors relate to one another and is used to develop rates, premiums, or rating factors. A third-party pricing model vendor is anyone other than the licensee itself who develops, maintains, or distributes that kind of model for use in ratemaking. One key distinction is that vendors do not assume any insurance risk from the model's output; they are supplying a tool, not underwriting risk themselves. Underwriting models work similarly but serve a different purpose. Rather than developing rates, an underwriting model relates risk characteristics and exposures to expected outcomes in order to inform or automate accept, decline, or modification decisions. A third-party underwriting model vendor is an outside party that builds or distributes that model. Again, the vendor itself carries no insurance risk from that output; the insurer using the guidance retains that responsibility. Pricing data refers to the facts or information from a source other than the licensee that either feeds directly into a pricing model or informs which inputs are selected, including output from another vendor's model, meaning model chains are captured by the definition. Underwriting data mirrors pricing data, but on the underwriting side, and refers to the information used as inputs to an underwriting

model or that informs which inputs are selected for evaluating accepting, declining, or modifying the terms of risks. These categories may overlap in practice in that a single entity might supply both data and models.

Lapham said the proposed framework rests on three pillars: 1) a voluntary, shared multi-state registry hosted by the NAIC, where vendors register and provide governance information once and multiple states can access that information rather than each state building its own separate system; 2) rate and underwriting filing requirements, which do not replace existing state filing processes but establish requirements around third-party data and models; and 3) ongoing requirements, namely an annual attestation in which a senior leader must certify that governance is effective, the data or models comply with applicable insurance law, and material and change notifications where vendors advise when they make significant modifications to a data set or model, including decommissioning that model.

Lapham elaborated that the proposed registry acts as a central repository for corporate governance information, model and data governance programs, and a list of insurers that have been supplied with a given data set or model. As a condition of registering, vendors would commit to responding to regulator requests within specified time frames, and the registry would indicate which type of registration applies in each state. He said vendors would be expected to document corporate governance information (contact information, legal entity and ownership structure, and responsible officers), model documentation (purpose, assumptions, inputs, limitations, performance metrics, and validation processes), and data documentation (accuracy, completeness, timeliness, auditable data lineage, quality controls, and a defined process for missing or incomplete data). Vendors must also disclose how they use consumer data and how consumers may access or correct their own records. He added that governance programs would need to align with the "NAIC AI Principles on Artificial Intelligence (AI)" (AI Principles) and the NAIC *Model Bulletin on the Use of Artificial Intelligence by Insurers*.

Regarding regulator access, Lapham said registered vendors could be asked to provide detailed materials, including model documentation, input and output specifications, validation and performance results, fairness and bias testing, and change logs and audit trails. He said the NAIC is not subject to open records laws and would treat registry submissions as confidential, but once information is shared with a state, that state's open records laws would govern the information. If a vendor designates information as confidential, the NAIC would commit to sharing it only with states that have the legal authority to maintain that confidentiality. Insurers must continue to file rates and underwriting rules in accordance with each state's requirements. This proposal would additionally require that third-party data or models used by an insurer also be filed, either as part of the insurer's own filing or as a separate filing submitted directly by the vendor. If a vendor does not cooperate, a state insurance regulator could prohibit insurers in that state from using the vendor's data or models. For states without authority to approve a third-party filing directly, the third-party filing would be folded into the insurer's own filing, with the insurer bearing the burden of demonstrating that the resulting rates and underwriting decisions are sound.

Lapham said vendors would be expected to update their governance documentation and contact information annually as needed. Additionally, a senior leader with relevant technical expertise and organizational authority would be required to attest annually that: 1) the vendor's governance program is implemented and effective; and 2) the vendor's data and models comply with applicable insurance laws in the states where they are used. Vendors would be required to submit material change notifications to the registry in a timely manner whenever they make a material modification to, or decommission, a model or data set so that state insurance regulators are not required to wait for the next annual attestation cycle to learn of such changes. He said regulators could then request further information if a change appeared likely alter a model's behavior a risk profile, or have a meaningful consumer impact.

Lapham emphasized that the registration, attestations, and all vendor-facing requirements exist alongside, not instead of, insurers' own obligation He said insurers must still assess whether a model is suitable for their own

purposes; conduct their own validation, testing, and monitoring; and ensure that any third-party data is complete, accurate, and appropriate. Insurers remain responsible for compliance with each state's applicable insurance code. They must include contract terms with third parties that provide the insurer and state insurance regulators with access to model and data information. None of these insurer-level obligations are removed because a vendor is registered.

Lapham summarized four key takeaways: 1) the current draft reflects phase one of a larger effort, tightly scoped to P/C pricing and underwriting, with other insurance functions potentially added in later phases; 2) the proposed structure rests on three pillars—a voluntary registry, filing requirements, and ongoing notification requirements (annual attestation and material change notifications); 3) insurer accountability is not altered, as the proposed vendor oversight in no way substitutes for an insurer's own validation, monitoring, and compliance obligations; and 4) the draft has already been shaped by extensive stakeholder input, including 23 comment letters received during the prior exposure period that generally pushed for a more risk-based, outcomes-focused approach.

Lapham concluded with a recap of the framework's status, which is an exposed draft at this point, still open for comments and revisions.

Chou commented that the current wording of the annual attestation provision is effective, which requires a senior leader with relevant technical expertise with former organizational authority to approve data and model governance attestations. He suggested that third parties should submit an annual plan for the number of AI system models anticipated to be implemented in the next year so that regulators can plan ahead with needed resources.

Guzman relayed a question from his chief actuary regarding whether it was the intention or whether it is implicit that one state would be able to see other states' responses to third-party filings in order to assess whether concerns raised by one state should be a concern to others, or whether each state's responses would be kept confidential. Lapham said he was not sure the framework contemplates that, but he agreed it was a question the Working Group should consider further.

Dahl commented that his experience is that tiering is usually pricing, while carrier placement can be part of underwriting. Lapham agreed. Stolyarov commented that Nevada considers tiering to be part of rating, including when it determines company placement.

Eric Ellsworth (Consumers' Checkbook/Center for the Study of Services) commented that he believes that for data vendors and/or model vendors, there will be a universe of companies that are already closely involved with the insurance industry and understand the kind of regulatory approach, but then there will be vendors that are further away that may not have any experience with this regulatory approach. Regarding the potentially large amount of personal information, he asked what might happen if data aggregators and vendors do not have any of this documentation or process controls. He asked whether the framework or insurance oversight addresses that gap. Lapham assumed the expectation is that such governance processes would be in place and that registration would hinge on vendors having those practices. However, he acknowledged the framework does not explicitly address the practical question of what happens if a vendor lacks them, and he said the Working Group would need to think through that process. Ellsworth further asked whether the framework contemplates integration testing between vendors and insurers and a documented fallback plan in the event a vendor's system or interface does not work, noting that accountability and integration issues often fall on the consumer when a problem crosses the boundary between third parties. Lapham requested that Ellsworth follow up with those concerns in writing.

Erica Eversman (Automotive Education & Policy Institute—AEPI) commented that she believes there will be a lot of vendors that are not familiar with these types of reporting requirements, and even though it is voluntary, she

asked what would encourage the vendors to enroll in the registry. Lapham said that incentives would likely come through the carriers with which vendors contract, as carriers may see an advantage in that registration. He said this was a worthwhile question and encouraged Eversman to submit written comments, noting the Working Group may need to further consider the issue and refine the framework. Eversman asked if registration is advantageous for both state insurance regulators and insurers, then why not make registration mandatory.

Scott Kosnoff (Faegre Drinker) asked whether NAIC's role in reviewing registrations would be a binary check of whether required documentation was included or would involve a qualitative assessment of the documentation, such as the adequacy of a vendor's governance program. Lapham said that is another process question that involves balancing states' responsibilities and their comfort level, with offloading review work to NAIC staff. He noted the NAIC Model Review team already provides technical guidance to states upon request but does not make final determinations regarding a particular state's requirements.

Anthony Habayeb (Monitaur) said the current draft reads more as a structured repository of information than as an active review process. He asked whether the framework envisions further clarifying that distinction, given the potential for conflict between information in the repository and information a state separately requests as part of its own filing review, and whether any centralized review is contemplated or whether the registry would instead operate similarly to how states already work, providing help only when a state asks for it. Lapham agreed with that characterization. He said that the Working Group would need to think through what that would look like in practice and welcomed written feedback.

James Mao (Grant Thornton) asked how the Working Group and the NAIC are factoring in state-level data broker and data collector laws, such as New Jersey's law (currently in effect but with delayed enforcement) and California's Delete Act and Delete Request and Opt-Out Platform (DROP), into the framework, given industry uncertainty about which regulatory regime would apply. Lapham asked for clarification on whether Mao meant state-level requirements affecting third parties outside the insurance context or those that would apply within the insurance context. Mao said both are possible, noting that data broker laws typically apply broadly across industries while the NAIC's focus is specific to insurance. Lapham said he did not have a ready answer but agreed the issue warranted further consideration by the Working Group, and he encouraged Mao to follow up in writing.

Lapham noted the relatively short turnaround between the July 8 exposure and the meeting. He encouraged Working Group members and interested parties to continue reviewing the draft over the coming weeks and to submit written comments in advance of continued discussion at the Summer National Meeting.

Having no further business, the Third-Party Data and Models (H) Working Group adjourned.

SharePoint/NAIC Support Staff Hub/Member Meetings/H Cmte/2026_Summer/WG-Third-Party/20260716 Interim-Meeting/Minutes-Third-PartyWG071626-Final.docx

