1. Consider Adoption of its Spring National Meeting Minutes

Attachment 1

–Commissioner Nathan Houdek (WI)

Draft: 3/31/26

Big Data and Artificial Intelligence (H) Working Group
San Diego, CA
March 24, 2026

The Big Data and Artificial Intelligence (H) Working Group of the Innovation, Cybersecurity, and Technology (H) Committee met in San Diego, CA, March 24, 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); Richard Fiore and Mark Fowler (AL); Jimmy Harris (AR); Lori Munn (AZ); Ken Allen (CA); Michael Conway and Jason Lapham (CO); Wanchin Chou (CT); Karima M. Woods (DC); Nicole Crockett (FL); Matt Kilgallen (GA); Weston Trexler (ID); Jack Engle (IL); Meggan Brumbaugh (IN); Shaun Orme and Shawn Boggs (KY); Caleb Huntington and Jackie Horigan (MA); Sandra Darby (ME); Kate Stojsih and Joe Stoddard (MI); Phil Vigliaturo (MN); Jo LeDuc (MO); Jacqueline Obusek and Angela Hatchell (NC); Colton Schulz and Matt Fischer (ND); Connie Van Slyke (NE); Christian Citarella (NH); Vanessa DeJesus (NM); Gennady Stolyrov II (NV); Avani Shah (NY); Matt Walsh (OH); Nicole Nash and Brian Downs (OK); Michael Humphreys (PA); Matt Gendron (RI); Andreea Savu (SC); Travis Jordan (SD); Emily Marsh (TN); Nicole Elliott and Amanda Crawford (TX); Eric Lowe and Michael Peterson (VA); and Lela D. Ladd (WY). Also participating was: Tregenza Roach (VI).

1. Adopted its Feb. 17 Minutes

The Working Group met Feb. 17 and took the following action: 1) adopted its Feb. 9 minutes; and 2) continued to discuss edits to the artificial intelligence (AI) systems evaluation tool and heard feedback from interested parties.

Block made a motion, seconded by Commissioner Ommen, to adopt the Working Group's Feb. 17 minutes (Attachment One-A). The motion passed unanimously.

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

Commissioner Houdek stated that the AI systems evaluation tool pilot process began earlier in March and that all participating states have sent, or are in the process of sending, inquiries to their respective insurance companies. The group of pilot state insurance regulators is meeting weekly to share insights on anticipated responses and to receive training on data science, compliance, and governance concepts reflected in the tool. The decision on which companies to include in the pilot is up to each pilot state's domestic regulator. Participating states are engaging with their domestic companies in advance of sending the pilot inquiry to answer questions and ensure that companies understand the goals of the pilot process. Where a company is part of a group, the pilot state regulators are coordinating efforts, including reaching out to non-pilot states to raise awareness that the company is included in the pilot. In general, there is a mix of companies across the main product lines, with more property/casualty (P/C) and life insurers than health insurers being selected. Most of the pilot states selected one to 10 insurers, with two states sending inquiries to more than 10 insurers.

Commissioner Houdek said the Working Group will first focus on Exhibit A to provide regulators with an overview of how companies use AI across their business operations. Then, states could use questions from Exhibits B-D to diver deeper into AI use. Pilot states are using the tool in support of a mix of regulatory processes (market conduct exams, financial exams, financial analyses, and as part of a more general regulatory inquiry). Pilot state regulators are collaborating on their insights to learn from each other. Throughout the pilot, the Working Group anticipates developing a coordinated mechanism to solicit feedback from participating companies on the tool. Recognizing

the broad interest in this initiative among regulators, industry, and interested parties, members of this Working Group and NAIC committee support have provided updates on the pilot at seven different meetings during the Spring National Meeting, including the three actuarial task force meetings, and the Market Regulation and Consumer Affairs (D) Committee, Financial Condition (E) Committee, and Innovation, Cybersecurity, and Technology (H) Committee. The Working Group will be providing public updates throughout the pilot process.

3. Received a Presentation on How to Operationalize the NAIC Model Bulletin on the Use of AI by Insurers

Dorothy Andrews (NAIC) presented possible approaches to move from adopting the *NAIC Model Bulletin on the Use of Artificial Intelligence Systems by Insurers* to actively implementing it through concrete governance, documentation, and oversight practices. She emphasized the importance of shared definitions for AI, machine learning, AI systems, consumer harm, and risk classes so that supervision is consistent and comparable across jurisdictions and companies. She stated that an AI risk taxonomy should serve as a guide to the appropriate level of governance, testing, and regulatory scrutiny for different AI uses. She outlined expectations for standardized AI governance that clearly describe executive accountability, board oversight, model inventories, data sources, and risk management frameworks. She highlighted the need for transparency around both internal and external data, including data purpose, demographics, bias and missing-data analysis, and third-party contractual and auditability and transparency considerations.

Andrews stated that AI model cards could be a practical tool to summarize what models do, how they should be used, how they perform, and their assigned risk level. She stressed the importance of ongoing monitoring to detect and respond to model drift rather than treating approval as a one-time event, and that attention should be given to protected class inference and bias testing, including the identification of proxy variables, statistical and sociotechnical analyses, and a clear discussion of model and data limitations. She noted that effective AI oversight requires sufficient human capital and expertise for both insurers and state regulators to review reports, follow up on risks, and enforce consequences where expectations are not met.

Commissioner Ommen, referring to the AI risk taxonomy of harm, commented that he has heard from interested parties that risk from the use of third-party data is mainly associated with pricing and underwriting, and asked Andrews to confirm. Andrews agreed, especially when it comes to third-party data that is aggregated but not sufficiently scrutinized by insurance companies, which could create harm consumers.

Eric Ellsworth (NAIC Consumer Representative) commented that two operational issues deserve attention. First is workflow integration, which refers to how an AI model is integrated into an insurance-related process. He stated that Andrews' presentation flags some impedance-mismatch issues with using a model outside its trained purpose, which are exacerbated when integrated into other processes. A component of that is the loss of knowledge by people who used to do the process manually. A strong governance structure raises awareness of knowledge loss and establishes a proactive stance to manage it. Second, there is a need for well-defined exception-handling processes when a model cannot do something well that it was not trained to do. Andrews replied that part of a good governance framework is ongoing training.

David Snyder (American Property Casualty Insurance Association—APCIA) reemphasized that one of the strengths of the NAIC AI model bulletin is that it is based on established regulatory standards and stated that it is important to constantly tether to those standards. He asked whether confidentiality will continue to be assured, given the assumption that the AI systems evaluation tool will be used in market conduct and financial examinations. Commissioner Humphreys responded that the Pennsylvania Insurance Department will administer the tool in a financial analysis context, thereby preserving confidentiality.

Commissioner Ommen echoed that the Iowa Insurance Division will be administering the tool through the financial examination process, where confidentiality would be preserved. In a market conduct review, Iowa would administer the tool under its market examination authority, and confidentiality would be maintained.

4. Heard a Panel Discussion on AI Governance Trends

Scott Kosnoff (Faegre Drinker) and Anthony Habayeb (Monitaur) discussed emerging best practices for AI governance, emphasizing that effective governance is primarily a business enabler rather than a narrow compliance exercise. The speakers highlighted the importance of cross-functional governance for aligning business, technology, risk, and compliance stakeholders and creating shared accountability. They stated the need for strong AI inventory management based on the principle that organizations cannot govern systems they cannot see. Success in many cases can be achieved by establishing awareness of the inventory of AI systems in place. The cross-functional dynamic is enforcing conversations and enabling proper governance and AI implementation success. They noted increasing pressure from insurance companies on vendors to demonstrate transparent, credible AI governance, especially for high-risk systems. The goal for insurers should be to have a clear story to tell regulators about their governance practices, demonstrating their appreciation for potential risks and their mitigation efforts. They described mature AI governance programs as those that institutionalize critical thinking about potential harms, impacts, and failure modes, rather than relying solely on formal processes.

Miguel Romero (NAIC) asked how to account for operational risk, financial risk, and the risk of adverse consumer outcomes, and whether companies are implementing a single risk criterion or a specific criterion for each type of risk. Habayeb responded that there is a learning curve, with many things that should be asked and managed upfront. Triage should be performed to determine the materiality of potential consequences that should be handled downstream. For example, if a system has been inherently evaluated as high risk, it should be assigned a control to address the risk of bias and its mitigation. Also consider robustness and system resiliency in performing risk assessment as a good fundamental risk management approach. Kosnoff stated that his firm has worked with clients to develop an AI use-case intake form to be submitted to the governance committee, which includes questions about the potential impact on policyholders, applicants, employees, other constituents, and the company itself in concern of a weakened financial condition. This intake form forces critical thinking about what the governance committee is asking to review, institutionalizing it and providing a starting point for the governance committee to conduct its review and determine where time and energy should be spent. Habayeb responded that the committee cannot reasonably review everything, so triage upfront is important for mitigation. Commissioner Houdek acknowledged that the tool's scope is intentionally broad at this point, but that it may be narrowed as a result of the pilot process.

Block asked the panel whether they are seeing any stratification by product line or company size in how well governance is progressing and maturing, and whether they are noticing any particular areas within companies that are driving the governance work. Kisnoff responded that he has not seen a stratification. Habayeb responded that large organizations tend to staff more full-time resources, but this may not be the case as much for regionals and smaller carriers, where the burden is larger. Kosnoff stated that it is hard to overemphasize the importance of education and training and stated that like the concern for model drift, there is a concern about human drift, when the people start to cut corners, get stressed, run short on time, get complacent, and become overly reliant on machines to make recommendations or determinations, which poses a compliance risk for all organizations.

Stolyrov II asked the panel for their insights on the extent to which a taxonomy of harm has become embedded in insurers' cultures, and whether to adopt a particular model. If an insurer could prioritize controls to prevent a

large amount of shock and outrage by consumers, prevent consumer complaints, regulatory actions, and litigation in the longer term, then that is a profitable series of moves, and ask whether they are seeing insurers think in these ways. Habayeb reinforced his earlier point that performing an intake assessment is an exercise in thinking about business impact.

5. Heard a Federal Update on AI

Shana Oppenheim (NAIC) said that on March 20, 2026, the White House released a National Policy Framework for Artificial Intelligence, outlining policy recommendations to guide Congress in developing a unified federal approach to AI legislation and regulation. The policy recommendations are consistent with what the Trump administration has been signaling about its views for some time: 1) the proliferation of state AI laws is creating barriers to innovation; 2) there needs to be some national standard governing AI; and 3) there are particular areas in which Congress should act in order to protect individuals from potential individual and economic harms that could be caused by the continued adoption of AI technologies.

Oppenheim added that last summer, the Trump administration urged Congress to adopt a temporary federal moratorium preempting certain state AI laws, but Congress ultimately declined to pursue that approach. Shortly thereafter, in December 2025, the administration issued Executive Order 14365, establishing a national policy framework for AI and seeking to curtail the impact and continued proliferation of state AI regulation. Among other elements, this Executive Order directed the Department of Justice to establish an AI litigation task force and instructed federal agencies to assess whether discretionary funding programs could be used to discourage certain types of state AI regulation. The Executive Order committed the administration to working with Congress to develop a federal legislative framework that would replace most state-level AI laws with a unified national standard. The framework released on March 20 follows up on that commitment by outlining the administration's preferred approach to federal AI legislation, providing guidance on the key areas that any federal legislation should address and the categories of state AI laws that should be preempted.

While the framework spans a wide range of policy areas organized around six key objectives, the following takeaways are important for companies that develop, augment, deploy, or test AI systems:

1) Child Safety and Privacy Regulation: A significant focus of the framework is protecting minors from AI harms and empowering parents to control their children's digital environments. The Framework encourages Congress to adopt age-assurance requirements for AI platforms likely to be accessed by minors, tools that can be used by parents and guardians to manage privacy and engagement settings, and limits on data collection and online behavioral advertising. The framework urges pursuing these goals while avoiding ambiguous content standards or open-ended liability regimes that could drive excessive litigation.
2) Copyright, Fair Use, and the Judiciary: The framework acknowledges the judiciary's authority to assess copyright and fair use questions related to AI training, while noting the administration's view that training AI models on copyrighted material does not violate copyright laws.
3) Antitrust Liability Exemption for Collective Licensing Negotiation: The framework encourages Congress to consider enabling licensing or collective-rights frameworks that would allow intellectual property rights holders to collectively negotiate compensation from AI model developers without incurring antitrust liability.
4) Free Speech Protection: The framework emphasizes limits on the federal government's authority to coerce AI providers to restrict or alter content for partisan or ideological reasons. It also directs Congress to provide avenues for redress where such coercion occurs.

5) No New Federal Rulemaking Body: The framework encourages relying on existing sector-specific regulators and industry-led standards rather than creating a new, centralized federal AI regulatory authority.
6) Federal Preemption of State AI Regulation: The framework supports broad federal preemption of state AI laws that impose undue burdens, while preserving states' traditional police powers to enforce laws of general applicability, especially to protect children, prevent fraud, and safeguard consumers. Additionally, the framework calls for precluding states from regulating AI model development or imposing liability on AI developers for unlawful conduct by third parties using their systems.

Oppenheim concluded that the framework is not a binding document and, by itself, does not impose new legal obligations or direct agencies to take specific regulatory actions. Instead, it outlines a series of recommended policy approaches for Congress to consider in drafting comprehensive federal AI legislation.

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

SharePoint/NAIC Support Staff Hub/Committees/H CMTE/2026_Spring/WG-BDAI/Spring-Minutes/Minutes-BDAIWG-032426 - Final.docx

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

–Coral Manning (WI)

Pilot Participation States

1. California
2. Colorado
3. Connecticut
4. Florida
5. Iowa
6. Louisiana
7. Maryland
8. Pennsylvania
9. Rhode Island
10. Vermont
11. Virginia
12. Wisconsin

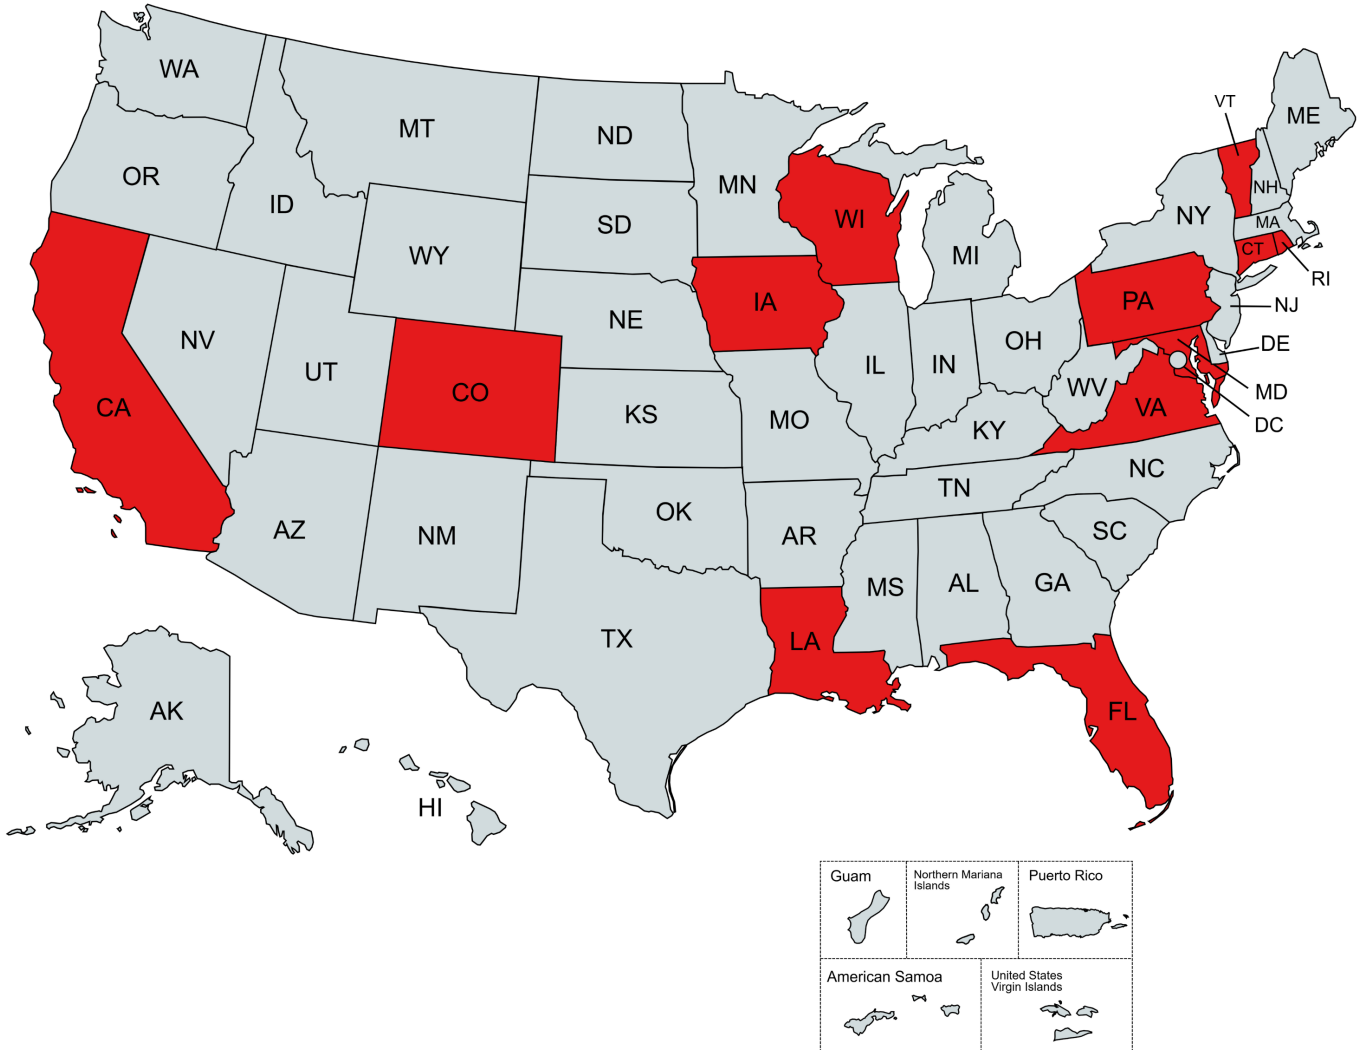

Update on the AI Systems Evaluation Tool Pilot Process

- Methods of implementation
  - Formal examination
  - Survey or data call
  - Other/hybrid

- What we're learning
  - Carriers are eager to provide feedback (survey developed)
  - Collaboration among states has been invaluable
  - Clarification is needed on some questions and definitions
  - There is a wide range of AI use across lines of business and carriers
  - How responses in earlier exhibits inform the use of later exhibits

AI Systems Evaluation Tool–Timeline

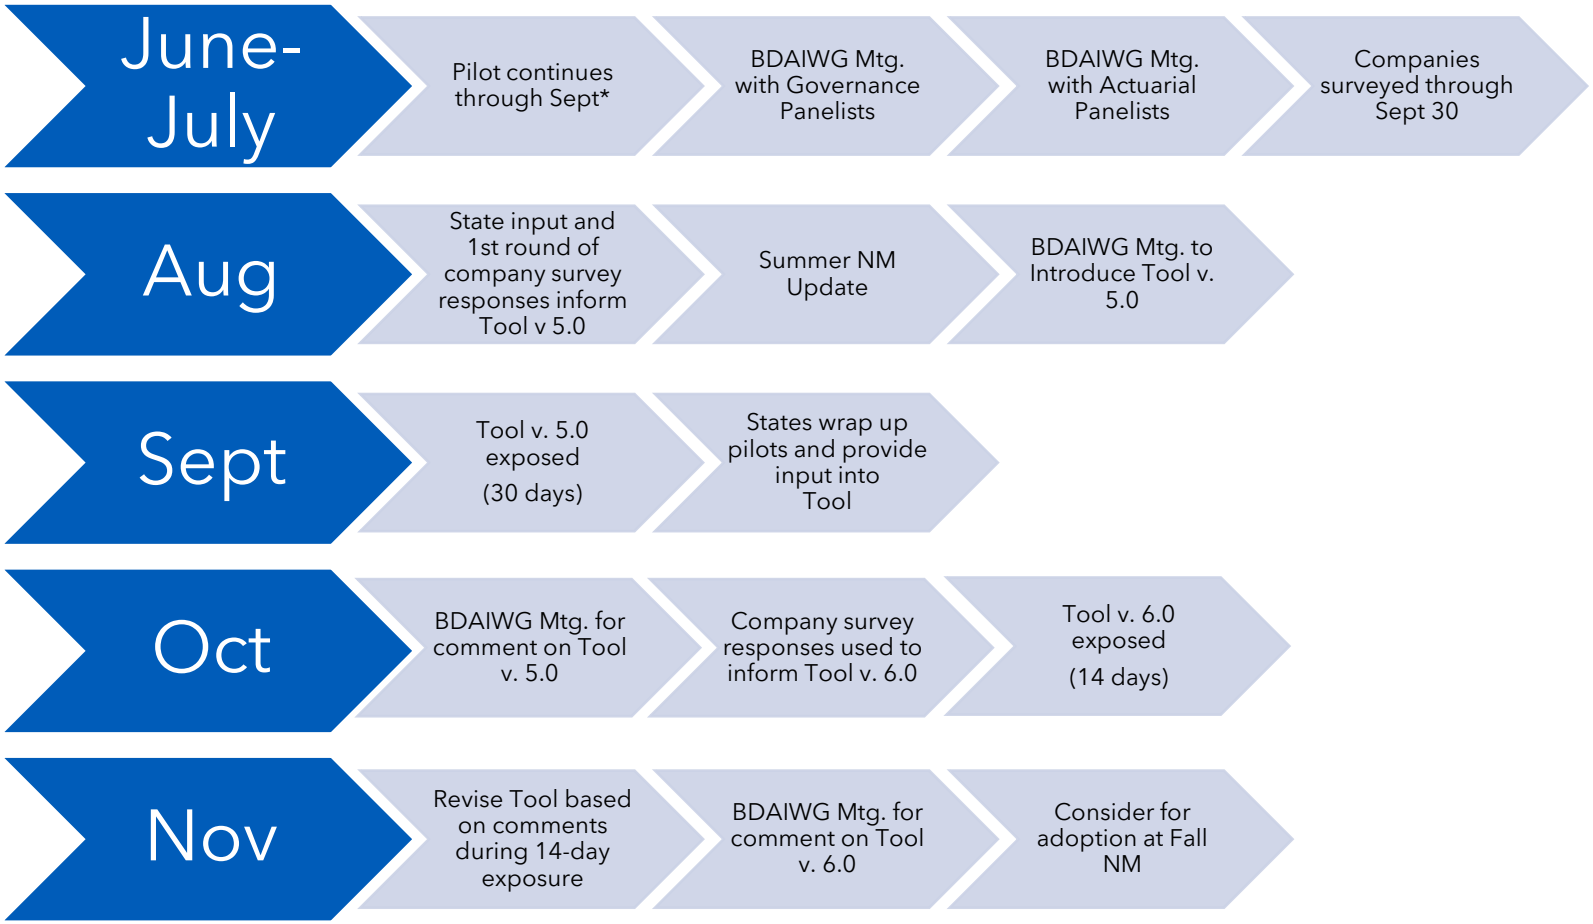

*Timing for each state's pilot will vary and every pilot may not be completed by 9/30/26. Information gleaned from the pilot experience will inform Tool drafts through adoption and future iterations of the Tool.

3. Hear a Panel Discussion on AI Governance Trends

–Michael David Ruiz (Deloitte)
–Fred Karlinsky (Greenberg Traurig, LLP)
–Mary Jane Wilson-Bilik (Eversheds Sutherland)