ASPEN POLICY ACADEMY
aspen institute

Model Framework

GenAI Forward Investigation Framework

By Michelle Sipics

March 2026

This policy project was developed by an Aspen Policy Academy fellow while participating in the Science and Technology Policy Fellowship. This document presents a framework, called GenAI Forward, for investigating generative artificial intelligence incidents. The full project, including a policy brief explaining the fellow's core recommendations, is available here. Please note that the author's opinions published here are their own. This publication does not reflect the views of the Aspen Policy Academy or the Aspen Institute.

The GenAI Forward investigation framework provides a set of guidelines for investigating certain types of harm-causing incidents involving generative artificial intelligence (GenAI) systems. Organizations can use this framework to incorporate best practices from safety-critical sectors into their investigations, with the goal of producing findings and recommendations that would increase the safety of GenAI applications. This framework was developed specifically for the Utah Office of Artificial Intelligence Policy, but it can be applied to other qualifying GenAI systems as well.

Overview

While GenAI continues to evolve rapidly – developing new capabilities, new applications, and potentially new harms – risks from emerging and continuously changing technologies are not new. Industries and sectors such as aviation, medicine, and cloud providers have long dealt with issues of safety; they have built trust and confidence in the safety and reliability of their products and services through their approaches to incident investigation and post-investigation recommendations. The GenAI Forward investigation framework leverages best practices and lessons learned from other sectors while acknowledging the scale and form of impacts that are specific to GenAI technology. The purpose of the framework is to enable continued safe innovation with GenAI and to reduce the occurrence and impact of GenAI incidents.

For the purposes of the GenAI Forward framework, a GenAI incident is defined as follows: **A GenAI incident is an event, circumstance, or series of events in which the development, use, or malfunction of a generative artificial intelligence system causes direct harm to an operator, user, or person(s) subject to decisions based on the outputs of the system.** For more about this definition and examples of relevant real-world incidents, see GenAI Incident Definitions and Examples.

How to Use this Framework

The GenAI Forward investigation framework is designed to be used for a root cause investigation – one that seeks to uncover all the reasons underlying the incident – after all immediate incident response activities have concluded. This framework assumes that the necessary actions have been taken to stop or address the harm specific to the incident; the goal is now to identify the factors that contributed to the incident.[1]

Organizations can use this framework to incorporate best practices from safety-critical sectors into their investigations of GenAI incidents. The GenAI Forward investigation framework comprises 6 steps:

1. Incident description, baseline information collection, and designation of investigation lead
2. Designation of investigation parties and participants
3. Discussion of investigation methodology
4. Designation of investigation groups and responsibilities
5. Designation of target questions
6. Discussion, analysis, reporting, and recommendations

The following sections discuss each step in turn.

Participants in the investigation should record the findings and outputs of each step. The process culminates in a shared discussion and analysis of the information gathered and, ultimately, enables users to create a report that makes recommendations for reducing the likelihood of and/or harm from incidents.

Note: Because incidents vary significantly in their severity, complexity of circumstances, and number of people and organizations involved, the time required to complete an investigation also varies widely. Therefore, this framework offers guidelines for investigation but does not provide a timeline.

1. Incident Description, Baseline Information Collection, and Designation of Investigation Lead

This step sets the scope of the investigation, documents the actions taken prior to the investigation (during incident response), and designates an investigation lead.

1a. Describe the Incident Under Investigation

The incident should be described in detail, specifically addressing the action and harm elements of the GenAI incident definition.

This common practice of documenting the incident – drawing from investigations in aviation, medicine, and other safety-critical sectors – helps create a common understanding at the beginning of the investigation. It also establishes a baseline for comparison at the end of the investigation. This is important because initial assumptions about the actions and factors that led to an incident are likely to be incomplete or incorrect.

For example, initial high-level review may suggest that problematic training data led to an incident (i.e., "development" is the relevant action from the definition). After the investigation has been completed, however, it may become clear that the way a model was incorporated into a workflow also contributed to the incident (i.e., both "development" and "use" are relevant actions). Capturing a baseline understanding of the incident before the start of the investigation creates a point of comparison for the post-investigation understanding. This helps underscore the value of the investigation process and provides a starting point for post-investigation recommendations. It also helps those who are unfamiliar with the investigation process become more comfortable with an evolving understanding of the incident — something that is common across investigations in all sectors.

1b. Collect Baseline Information

All actions taken during the initial response to the incident should be collected and documented, to be shared with participants in the investigation along with the incident description.

1c. Assign an Investigation Lead

The designated investigation lead should be a representative of the organization that has incorporated the GenAI model involved in the incident into its employee, customer,

or user workflow or engagement. For example, if Company X is using a GenAI model to customer service requests, and that model was involved in an incident, a representative from Company X should lead the investigation (regardless of who created the GenAI model). The investigation lead may be selected based on the context of the incident, prior incident investigation experience, or both. Designating a single lead is common practice in incident investigation across sectors. Doing so creates a universal source of information, communication, and coordination throughout the investigation.

The designated lead will be responsible for the following:

- Coordinating communications and meetings between the various parties to the investigation;
- Following up with representatives of parties to the investigation to ensure that requested information is provided in the level of detail and format requested, in accordance with investigation timelines;
- Responding to questions from group representatives and participating experts;
- Assembly of the final investigation findings; and
- Dissemination of the final investigation findings to all parties.

2. Designation of Investigation Parties and Participants

This step establishes the structure of the investigation's parties and participants.

2a. Determine the Relevant Parties to the Investigation

First, the relevant parties to the investigation must be determined. The party system is a hallmark of aviation incident investigations performed by the National Transportation Safety Board (NTSB), which are widely considered to be the gold standard in investigating transportation accidents and incidents. According to the NTSB's Party Guidance:

> *"When the NTSB launches to an investigation, we assemble a broad spectrum of technical expertise to investigate complex transportation accidents. We designate as parties to the investigation entities whose employees, functions, activities, or products were involved in the accident or incident. This facilitates the rapid and complete acquisition of all relevant factual information. Only those entities that can provide technical expertise or knowledge to an NTSB investigation are granted party status, and only those persons who can provide the NTSB with needed technical expertise or specialized knowledge are permitted to participate in an investigation."*

Jim Hall, a former chair of the NTSB, explains further: "The underlying premise of the party system is a strong one: that everyone has an overriding interest in safety and that everyone wants to find out what happened so that steps can be taken to ensure such an accident is not repeated."[2]

In an aviation incident investigation, for example, the parties are likely to include the aircraft manufacturer, representatives of the pilots or crew (such as the Air Line Pilots Association), the Federal Aviation Administration (which may also represent any air traffic controllers involved in the incident), the manufacturer of the aircraft's engines, and so on. Which parties are invited to participate in the investigation depends on the nature of the incident.

GenAI incidents may or may not involve a physical outcome the way that a plane crash does, but they still involve multiple parties. The following are some of the potential parties to a GenAI incident and their representatives:

- **The organization that is using GenAI in its operations, services, or products.** The investigation lead will come from this organization, but they are unlikely to be the only person from this organization who should participate.
- **The organization that created the GenAI model.** In the case of ChatGPT, the representative might be a technical representative from OpenAI; for Copilot, from Microsoft; or for Q, from Amazon. If the organization using the GenAI model built the model in-house, then they are one and the same.
- **The organization that trained the GenAI model.** This may be the same as the organization that created the model, but the representative(s) of this party should have expert-level understanding of the training data and the methodology used to train it.
- **The organization that implemented the GenAI model for the organization using it.** In the case of models built in-house, this may or may not be the same organization that created the model. Otherwise, it is likely to be a technical representative from the "host" organization that has incorporated the model into an employee workflow, user offering, or customer experience, or a representative from a third-party firm hired to perform the implementation work.
- **The organization that operates the platform on which the model was being used at the time of the incident.** This may be the same as the other parties listed above (e.g., a chatbot incorporated into a company's website, a standalone app), but GenAI models can be made available to users or operators via other platforms. For

example, a GenAI model accessible via chat may be made available to employees or customers via Slack, Discord, Microsoft Teams, and so on. If the incident occurred while the user or operator was accessing the model via such a platform, a representative from the platform should be appointed to the investigation.

- **The organization that designed and operates the user interface through which the model was accessed.** This may be the same as the parties outlined earlier, or it may be a secondary organization. In either case, the party should appoint a representative who has an understanding of the interface that the user leveraged during the incident.
- **The organization that built the physical device on which the end user accessed the GenAI model.** For example, if the incident in question involved a user accessing a voice-driven GenAI model via a smartphone, a representative from the smartphone manufacturer can be appointed.

The investigation lead should use their discretion to ensure that, as in the NTSB Party Guidance, entities whose employees, functions, activities, or products were involved are represented as parties to the investigation.

2b. Determine Additional Participants Whose Experience or Expertise Is Relevant to the Investigation

The investigation lead should next determine at least one designated point of contact or individual for each category of additional relevant experience or expertise. These participants may be interviewed by the investigation lead and participate in investigation groups (see step 4) to create a more complete understanding of the incident. These categories may include, but are not limited to, the following:

- **System user(s) or operators.** In any GenAI incident being investigated using this framework, at least one human user or operator has interacted with the system in some way, whether by providing an input or prompting the system for some action, taking the output of the system to complete a larger workflow, or both. Depending on the system, this may be an employee, a customer, or another end user. This person should be interviewed extensively to document their recollection of the incident, preferably as soon as possible after it occurs. If the person(s) was interviewed during incident response, they should be interviewed again during the investigation phase. The interview should document their familiarity with the system; the way(s) in which they typically access the system (both during the incident and at other times); their goals for the use of the system; any incentives or disincentives

around their use of the system; and (to the extent possible) their state of mind before, during, and following the incident.

- **Domain experts.** GenAI systems, like all technologies, are not applied in a vacuum. If a model is being used in a customer service capacity, for example, an expert in customer support or engagement should be given the details of the incident and interviewed to provide relevant domain expertise on the use and implementation of the model within that context. This would also apply to industry-specific applications: For example, a model being used in a healthcare environment would require interviews with experts in the healthcare domain. Ideally, this would include expertise from both the organization that has incorporated the incident-involved GenAI model into an employee, customer, or user workflow or engagement, as well as external expertise, such as a representative from an industry group(s) or researcher(s) in the relevant space.
- **User experience experts.** Depending on the nature of the incident, it may be helpful to interview experts in user experience to understand how the larger circumstances surrounding the incident did or did not affect the employee's, customer's, or user's interactions with the system.

3. Discussion of Investigation Methodology

This step creates a shared understanding of the approach to investigation.

The GenAI Forward investigation framework is based on the "just culture" concept, which has been applied in other safety-critical settings such as aviation and healthcare. This approach to systems thinking avoids a "who caused this problem" mode of inquiry in favor of asking the more open-ended "what went wrong?"

This approach still allows for individual and group accountability in the case of intentional misconduct or negligent behavior. But for incidents that are the result of honest mistakes or oversights – situations that are likely to occur when leveraging emerging technologies for which best practices may be evolving – this approach encourages open, honest communication about the events and decisions leading up to the incident. This improves the investigation's chances of uncovering the root causes of the problem so that they can be corrected, thereby improving the safety of the system and reducing the likelihood of future harm.

The investigation lead should explain this approach to each party's representative and to each participant in the investigation and ensure that they are willing and able to meet

the responsibilities of their role.

4. Designation of Investigation Groups and Responsibilities

This step designates investigation groups and outlines their responsibilities.

When investigating aviation incidents, the NTSB also specifies investigation groups, which might include the following:

- **Operations** — providing the history of the flight(s) involved in the incident as well as an understanding of crew members' work duties leading up to the incident.
- **Systems** — looking at the components of the aircraft's systems.
- **Air traffic control** — re-creating and reviewing the services provided by air traffic control before and during the incident.
- **Weather** — looking at all of the relevant weather data related to the incident.
- **Human performance** — studying the crew's performance during the incident and any factors that occurred before the incident that might have contributed to human error (e.g., fatigue, medication, medical histories, training, workload, equipment design, work environment).

This list is not exhaustive, and the existence of a group does not suggest that it necessarily had any involvement in the incident under investigation. For example, all NTSB investigations gather weather information, but many aviation incidents occur without weather contributing to the incident. Considering (and eliminating or confirming) the possibility, however, is still important.

The above groups illustrate the recommended approach to investigation: looking at as many factors as possible that may have contributed to the incident and creating a group that consists of experts in that space, both representatives of the organizations involved in the incident (parties) and those whose knowledge can help create a more complete understanding of the circumstances (participants).

Designating GenAI Investigation Groups

Examples of groups that could be relevant for GenAI incident investigations, and the responsibilities of each group, are outlined below. This list is not intended to be exhaustive. The investigation lead should determine which groups are necessary to complete a thorough root cause, just-culture-driven investigation.

Group	Responsibilities
Design group	Provide documentation and respond to questions about the development and design of the GenAI model under investigation, including any customizations of the application that are relevant to the incident. If available, provide any data retained related to the input, operation, or output of the model leading up to, during, and following the incident. This group should provide context and analysis related to the design and use of the model under investigation and how it may have contributed to or affected the incident outcome.
Implementation group	Provide documentation and respond to questions about how the model was implemented as part of an engagement or workflow, including any external information on which the model depends aside from user input, and how the model's output is used. If available, provide any data retained from the engagement or workflow leading up to, during, and following the incident. This group should provide context and analysis related to the implementation and workflow incorporation of the model under investigation and how it may have contributed to or affected the incident outcome.
Training group	Provide documentation and respond to questions about the data used to train the model and the methodology used for training, including ongoing training (e.g., does the model learn continuously and evolve its operation based on user interactions, including interactions during the incident?). This group should provide context and analysis related to the training of the model under investigation and how it may have contributed to or affected the incident outcome. *Note: In this case, the Training group refers to training for the GenAI model. A separate Employee Training group is listed later in this table.*

Platform group	Provide documentation and respond to questions about the platform environment in which the model was accessed during the incident. This may include real or exemplar screenshots, screen recordings, audio recordings, or other data representing the user's interaction with the model before, during, and following the incident (depending on the mode of engagement). It may also include other data from the user's time on the platform before and after the incident separate from their interactions with the GenAI model (e.g., active time spent on the platform before and after engaging with the model, other activities performed before and after the user's interaction with the model, errors encountered by the user during the larger incident session). This group may also consider whether the user accessed the model during the incident via a different platform than they typically use to interact with the model, which could also be a contributing factor to the outcome. This group should provide context and analysis related to the platform used to access the model and how it may have contributed to or affected the incident outcome.
User Interface group	Provide documentation and respond to questions about the specific interface from which the user accessed the GenAI model during the incident. This group may also coordinate with other groups to help determine whether the user interface used during the incident was different from that previously leveraged by the same user (e.g., if the user typically interacted with the GenAI model by typing questions but instead used spoken language interaction during the incident). This group should provide context and analysis related to the interface the user leveraged to interact with the model and how it may have contributed to or affected the incident outcome. *Note: There may be significant overlap between the Platform and User Interface groups.*

Device group	Provide documentation and respond to questions about the device from which the user accessed the GenAI model during the incident. This group may also coordinate with other groups (Platform and User Interface, in particular) to create a larger picture of the user's activity during the incident both related and unrelated to the model interaction. This group should provide context and analysis related to the capabilities of the device used to access the model and how those capabilities may have contributed to or affected the incident outcome. For example, if the user was leveraging a device's microphone to provide spoken input to a model, the fidelity of the microphone may have contributed to the model's "understanding" of the user's intentions.
User Understanding group	Provide documentation (if available), context, and analysis about the user's familiarity with the GenAI model in question, including their understanding of its capabilities, intended use, and how to interact with it to achieve the best outcomes. Representatives from multiple parties are likely to be involved in this group, which should also involve interviews with the user.
Employee Training group	In the case of incidents that occur when an employee is using a GenAI model as part of their expected workflow, investigators may consider including an Employee Training group to examine the training that the employee was given related to the use of the model in that workflow. This group can provide insights and analysis around the content of training, the manner and frequency in which it is delivered, and how those factors may or may not have contributed to the incident in question. Even if training is not found to be a contributing factor to the incident under investigation, investigators may recommend additional training based on findings from other groups to reduce the likelihood of incidents from other causes.

Domain group	Provide insights and analysis about domain-specific information and factors relating to the incident under investigation. In the case of a GenAI incident in a medical environment, for example, medicine would be the relevant domain. It is unlikely that the experts representing technical parties will have a thorough understanding of the domain to which their models, data, or devices are being applied. Bringing in expertise from the domain is critical to understanding whether GenAI models are being applied in a way that is safe based on the specific needs and conditions of that domain.

5. Designation of Target Questions

This step outlines the target questions that the investigation should seek to answer.

The National Institute of Standards and Technology (NIST), an agency within the US Department of Commerce, provides standards and recommendations for a variety of technology-related activities. NIST's Information Technology Laboratory publishes a widely adopted methodology for cybersecurity incident response. An earlier version of this methodology included a "lessons learned" section that listed questions to be addressed after an incident. Although these questions were written specifically for cybersecurity incidents, they can serve as a valuable set of guideposts for GenAI incident investigation.

The GenAI Forward framework uses a related but evolved set of questions to guide an investigation. This does not necessarily mean that asking these questions directly is the best course of action. Rather, parties to the investigation should gather information and conduct interviews with the ultimate goal of providing answers to these questions. Each investigation group should provide all data, insights, analysis, and further questions relevant to their responsibilities to help provide a thorough answer to each guiding question.

The following table lists the guiding questions to be addressed by the investigation. For each question, examples are provided to illustrate the types of data that can be collected and/or questions that can be posed to investigation groups or other participants based on their responsibilities, experience, or expertise relevant to the investigation. This is not intended to be an exhaustive list of data to collect or avenues to explore. The investigation lead and the participating groups should use their

experience and expertise to shape the inquiries, data collection, and analysis toward the goal of answering the guiding questions.

Guiding Questions

Guiding Question	Examples of Data or Information to Collect and Questions to Ask
Exactly what happened, with time stamps provided for relevant items before, during, and following the incident?	• Logs of model inputs and outputs • What does the user remember about the GenAI interaction? • What was the user trying to accomplish while using the model? • What impact did the incident have on the user or operator?
Did the system perform according to the organization's policies for its operation during the incident?	• Logs of GenAI model inputs, outputs, and operations, with reasoning steps if captured • Organizational policies for GenAI model operation and use • Technical limitations in place to govern the model's operation
Did the user of the system follow instructions or policies for using the model or its outputs?	• System instructions and policies provided to users of the system ◦ May include prominent instructions as well as terms of service • What does the user remember about instructions related to the use of the system and its outputs?
Had any recent changes been made to the model or the larger system prior to the incident?	• System documentation • Changelogs and code commits • What were the goals of the changes made to the system? (posed to the Implementation group and other groups as relevant)

Were there any indications that an incident like this was possible that were missed before the incident?	• Prior system logs and outputs • Prior user reports • Had the user ever interacted with the system before the incident or had an experience that was similar in any way?
What information did the organization lack that could have prevented the incident or mitigated the harm from it?	• First report(s) of incident and all related communications • Relevant logs and when and how they were available, and to whom
What precursors or indicators could be used to detect similar incidents in the future?	• Should the input provided to the model during the incident be flagged for future monitoring? • What elements of the output generated by the model could be flagged for future monitoring?
What corrective actions could prevent similar incidents in the future?	• What technical solutions could be put in place to stop similar incidents? • Could a policy change regarding the model's use be made to prevent similar incidents?
What corrective actions could mitigate the harm from similar incidents in the future?	• What technical solutions could be put in place to limit or eliminate the harm from similar future incidents? • Could a policy change regarding the model's use be made that could limit or eliminate the harm from similar future incidents?
What additional tools or resources are needed to detect, analyze, and mitigate future incidents?	• Could more data be securely captured during system operation in case it is needed to detect and/or investigate future incidents? • Are more resources required to monitor system outputs for potential issues? • Is training required for system designers, trainers, maintainers, operators, or users to mitigate future incidents?

The investigation lead should interview the additional parties whose expertise has been deemed relevant to the investigation. At the investigation lead's discretion, other parties to the investigation may be invited to participate in these interviews and ask follow-up questions or provide additional context as necessary.

6. Discussion, Analysis, Reporting, and Recommendations

This step sets the approach for discussion, analysis, and creation of the final investigation report.

The target date of completion and intermediate deadlines for the investigation are left to the discretion of the investigation lead. However, the lead should share all planned deadlines and the specifics of what is due at each deadline for all investigation groups, taking into consideration the availability of relevant experts and/or users or operators for interviews.

6a. Meeting and Discussion

While the initial information-gathering phase of the investigation can be accomplished by disparate groups working separately, incidents are typically the result of multiple factors that influence one another.

Here is a hypothetical example: Imagine a user for whom English is not their first language, and they are more fluent in spoken English than in written English. They have spent a lot of time using a particular GenAI model, and they have found that it provides more useful responses when they give their input via spoken natural language instead of typing a formal written prompt. But today, they are accessing the model while trying to get to work on a windy day, and their smartphone's microphone cannot handle the wind noise. Instead, they try to type their input while hurrying down the street to catch their bus. The input that is ultimately sent to the model has a slight mistake that changes the meaning of the prompt, and the model understands a different input than the user intends. Unfortunately, the user is in a hurry and does not notice, and they make a decision based on an output that is based on the wrong input.

There are multiple factors at play here: the user's familiarity with the GenAI model, their preferences in using it, and the circumstances in which they were accessing it. Change any of those elements, and the problem might not occur. But if the groups looking at those factors do not discuss them together, they may never get a full understanding of why the problem happened.[3]

For this reason, investigations in nearly all safety-critical sectors involve one or more meetings in which the groups share their findings and participate in a larger discussion to build a comprehensive understanding of the incident.

The investigation lead should plan for at least one meeting of all the investigation groups after each group has made its contributions. All the groups are encouraged to submit proposed findings and conclusions that they believe should be drawn from the evidence obtained during the investigation. Groups may also propose recommendations to prevent future incidents.

Before the meeting, the investigation lead should compile the findings and recommendations of all the groups in a document and share it for review. All the groups should use this document to prepare for the meeting.

Each group should present during the meeting its relevant findings for discussion among the full group and answer any questions from other group representatives. Following the meeting, the investigation lead should prepare a draft investigation report that summarizes the conclusions of the investigation (using the guiding questions as a framework) for review and further discussion by the groups. If needed, additional meeting(s) can be scheduled at the discretion of the investigation lead for discussion, additional information collection, and analysis, until:

- All group representatives have no further questions; or
- The investigation lead determines there are no additional available sources of information to answer outstanding questions; or
- The investigation lead determines that exploring additional questions is unlikely to have a material impact on the investigation's findings.

6b. Final Report

A final draft of the investigation report should be prepared, including both the pre-investigation incident description and any adjusted description of the incident based on the investigation's findings. Following the description, the report should list all the factors that contributed to the incident. This list is not to be used for the purpose of assigning blame (except in the case of negligence or intentional misconduct). Rather, following the description of these root causes, the report should include recommended actions that can be taken by any involved party to mitigate future incidents.

These can include, but are not limited to:

- Training changes for existing or future models;
- Changes to the way the model is included in a workflow;
- Changes to the model's user interface;
- Changes to the availability of the model via certain platforms or devices;
- Changes to post-processing or analysis of a model's output before it is passed to the next place in its existing workflow; and
- Changes to the retention of information associated with use of the model (e.g., user input and resulting outputs).

Finally, the report should document what information would have been helpful in the course of the investigation that was not available or not made available. Involved parties should use these notes to determine what actions they can take to gather information ahead of future incidents or in response to future incidents (e.g., usage logs, encrypted logs of model inputs and outputs, earlier interviews with involved people following an incident report).

All the information gathered and analysis done throughout the investigation should be included as supporting material if practical (e.g., as appendices) or retained securely and made available only as necessary if required for security or privacy reasons.

6c. Endorsement of Investigation Report

All parties must review the final draft of the investigation report and submit their decision to:

- Endorse the report;
- Endorse the report with additional comments; or
- Refuse to endorse the report with comments.

It is left to the discretion of the lead investigator to finalize the report if a significant number of group representatives refuse to endorse it; however, all refusals to endorse must be noted in the final report, as well as any comments submitted with either an endorsement or a refusal to endorse.

If you'd like to learn more, see the full project, including a policy brief explaining the fellow's core recommendations, at aspenpolicyacademy.org/project/genai-forward-incidents-2026.

Endnotes

1. Actions taken and information documented during the incident response phase should be made available to those running the post-incident investigation. Organizations should thoroughly document all stages of incident response and make that documentation part of the official GenAI incident response guidelines, if that is not already a requirement. This is common practice in cybersecurity incident response and investigation and provides a helpful starting point for GenAI incident investigation — even for events that did not involve a security breach or malicious intent.

2. James M. Walters and Robert L. Sumwalt III, *Aircraft Accident Analysis: Final Reports* (McGraw Hill, 2000), xxviii.

3. We use the word "problem" rather than "incident" here because this hypothetical situation is too vague to fully meet our definition of a GenAI incident. However, we include it to illustrate how different factors can contribute to an unintended outcome.

ASPEN POLICY ACADEMY
aspen institute

aspen institute

About the Aspen Policy Academy

The Aspen Institute's Policy Academy helps community leaders and experts across the political spectrum elevate their voices, influence key decisions, and strengthen democracy from the ground up. Our innovative training programs and resources equip people across sectors – from tech to the environment, science to civic engagement – with the skills to shape critical policy efforts. Learn more at aspenpolicyacademy.org.

Photo credit: RDGraphics via Canva