

Informed Innovation: Establishing a Generative AI Incident Investigation Framework in Utah

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 policy brief provides an overview of the fellow's proposed solution. The full project 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.

Executive Summary

The Utah Office of Artificial Intelligence Policy (OAIP) should adopt a public framework for investigating harmful generative artificial intelligence (GenAI) incidents to encourage safe and continued innovation from organizations using this technology. The OAIP currently offers an optional “regulatory sandbox” program that gives companies a path to experiment with GenAI solutions and time to resolve any issues that arise as long as they remain in compliance with the program’s terms. The program currently puts guidelines in place for the adoption and use of GenAI, leveraging [evaluation criteria](#) proposed by a prior team of Aspen Policy Academy fellows. The OAIP should expand its guidance to establish a [GenAI Forward Framework](#) outlining how companies should investigate GenAI incidents that result in harm. The framework establishes a process that OAIP and partners can use to investigate harmful GenAI incidents, including encouraging organizations to publish investigation reports through the OAIP to increase transparency and shared knowledge in the field.

Background

The State of Utah [established the OAIP](#) in 2024 to study advancements in AI technology, support innovation by Utah AI companies, identify risks to consumers, and issue policy recommendations for the responsible use of AI. One mechanism that the office uses to deliver on its mission is a [“regulatory sandbox.”](#)

In this optional program, companies enter into an agreement that sets guidelines for the operation of GenAI technology and, as long as they remain compliant, they receive regulatory relief (a chance to address problems before regulators take action) during an agreed-upon period of GenAI exploration. This approach allows Utah to protect its residents while still encouraging innovation and growth.¹

The rapid evolution of GenAI technology and its applications presents constant challenges for both the OAIP and participating companies. Even with carefully developed guidelines in place, GenAI incidents are bound to occur, potentially causing harm and undermining trust and confidence in the technology and in the OAIP's credibility.

Research highlights many ways that existing GenAI systems could cause harm. There are also documented cases of real-world GenAI systems causing or creating potential for real harm, such as:

- Unfair decisions impacting finances or livelihood, such as biased hiring, loan approvals or rejections, or bad financial advice from models trained on biased data.
- Repercussions from illegal activity, such as someone following dangerous advice from a hallucinating GenAI chatbot.
- Physical or psychological harm, such as someone following harmful advice from a chatbot.
- Business harm or financial loss from a GenAI model bypassing safeguards.

Utah's existing guidelines are intended to mitigate harm by providing a responsible framework for the initial adoption of GenAI. The current regulatory sandbox agreement, however, lacks guidance on how to conduct an investigation when a harmful GenAI incident occurs.

Recommendations

The OAIP should expand the scope of its efforts to promote GenAI innovation and public safety by adopting a standardized framework for investigating GenAI incidents.

First, the OAIP's regulatory relief program should adopt the new GenAI Forward framework for incident investigation based on best practices and lessons learned from other sectors, such as commercial aviation and medicine. The GenAI Forward framework comprises 6 steps:

- Incident description, baseline information collection, and designation of investigation lead;
- Designation of investigation parties and participants;
- Discussion of investigation methodology;
- Designation of investigation groups and responsibilities;
- Designation of target questions; and
- Discussion, analysis, reporting, and recommendations.

This framework would encourage transparency and open discussion when GenAI incidents occur and support two goals: to prevent future harm and to allow companies to continue to innovate using GenAI technology. A common framework would strengthen the value of the OAIP's agreements, giving companies clear guidance and a means to continue to explore potential GenAI solutions.

The OAIP should share this framework, as well as relevant investigation results, with the public. "Showing the work" would increase public trust in GenAI and the OAIP as a leader in this space. It would also encourage widespread adoption of the GenAI Forward framework, thereby strengthening the larger ecosystem in which organizations operate and encouraging fair competition.

Require organizations participating in the OAIP regulatory sandbox program to adopt the GenAI Forward investigation framework

The OAIP should adopt the GenAI Forward investigation framework to help reduce the chance of, and harm from, future GenAI incidents. There has been significant discussion of responsible deployment of GenAI, but far less discussion of how to investigate what led to an incident and how to use those insights to prevent future incidents.² The OAIP could improve the value of its regulatory sandbox agreements for signees and mitigate the harm of GenAI incidents to Utah residents by adopting a standardized investigation framework.

Such a framework would not require starting from scratch. While there are many potential harms from GenAI applications, the risks from emerging and continuously evolving technologies are not new. Industries and sectors such as aviation, medicine, and cloud providers have long dealt with safety concerns and built trust and confidence in their solutions through their approaches to incident investigation. The proposed framework for investigating GenAI incidents incorporates best practices and lessons learned from other sectors in which trust and safety are critical, while also acknowledging the particular scale and form of impacts specific to GenAI technology.

Publish the GenAI Forward investigation framework and voluntary pledge

Incorporating the investigation framework into the OAIP’s regulatory relief agreements would help improve safety and outcomes for Utah’s residents while giving companies the opportunity to continue innovating with GenAI. To encourage safety more broadly and to increase general trust in its operations, the OAIP should share this framework publicly and encourage companies to sign a voluntary pledge (see Appendix). Organizations that sign the pledge would commit to using the framework in their own GenAI incident investigations and, whenever possible, sharing the findings, recommendations, and outcomes of relevant investigations. The pledge could be open to any organization and not limited to those with a presence or operations in Utah.

“Showing the work” would help increase public trust in GenAI and in the OAIP as a leader in this space. The increased disclosure about the OAIP’s operations would also underscore the agency’s value to Utah residents and consumers.³ Publicly sharing incident investigations would help highlight areas where the existing framework could be expanded or improved, so that it can continue to evolve as needed alongside the technology itself.

To recognize organizations that are working to improve the safety of their GenAI applications and to encourage continued adoption of the framework, the OAIP should create a webpage listing organizations that have signed the pledge and adopted the GenAI Forward investigation framework and feature those that have shared findings from their own investigations.

Endnotes

1. Companies that are actively participating in the OAIP’s regulatory relief program have touted its value. ElizaChat, for example, is a company that works to bridge gaps in traditional mental health support using GenAI. [ElizaChat CEO Dave Barney explained](#), “The AI Policy team engaged with us, understood our business needs, and crafted a regulatory relief contract that freed us to explore creative products that will help teenagers improve their mental health, without fear of regulatory risk.”
2. When interviewed, leaders at one of the companies participating in the OAIP’s regulatory sandbox said that a common framework for investigating GenAI incidents would be a valuable addition to the program.
3. In interviews conducted in July 2025, most Utah residents contacted were unaware of the OAIP and its mission.

Appendix: GenAI Forward Pledge

The following is a sample pledge for organizations and companies to sign.

The GenAI Forward pledge is a commitment to adhere to a best practices framework when investigating incidents involving generative artificial intelligence (GenAI). It is designed to unite organizations across geographic boundaries, industries, and missions in the responsible development, application, and evolution of GenAI. Organizations signing the GenAI Forward pledge agree to:



Adopt a common definition

The framework provides a best practices approach for investigating GenAI incidents that meet an agreed-upon definition: 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.

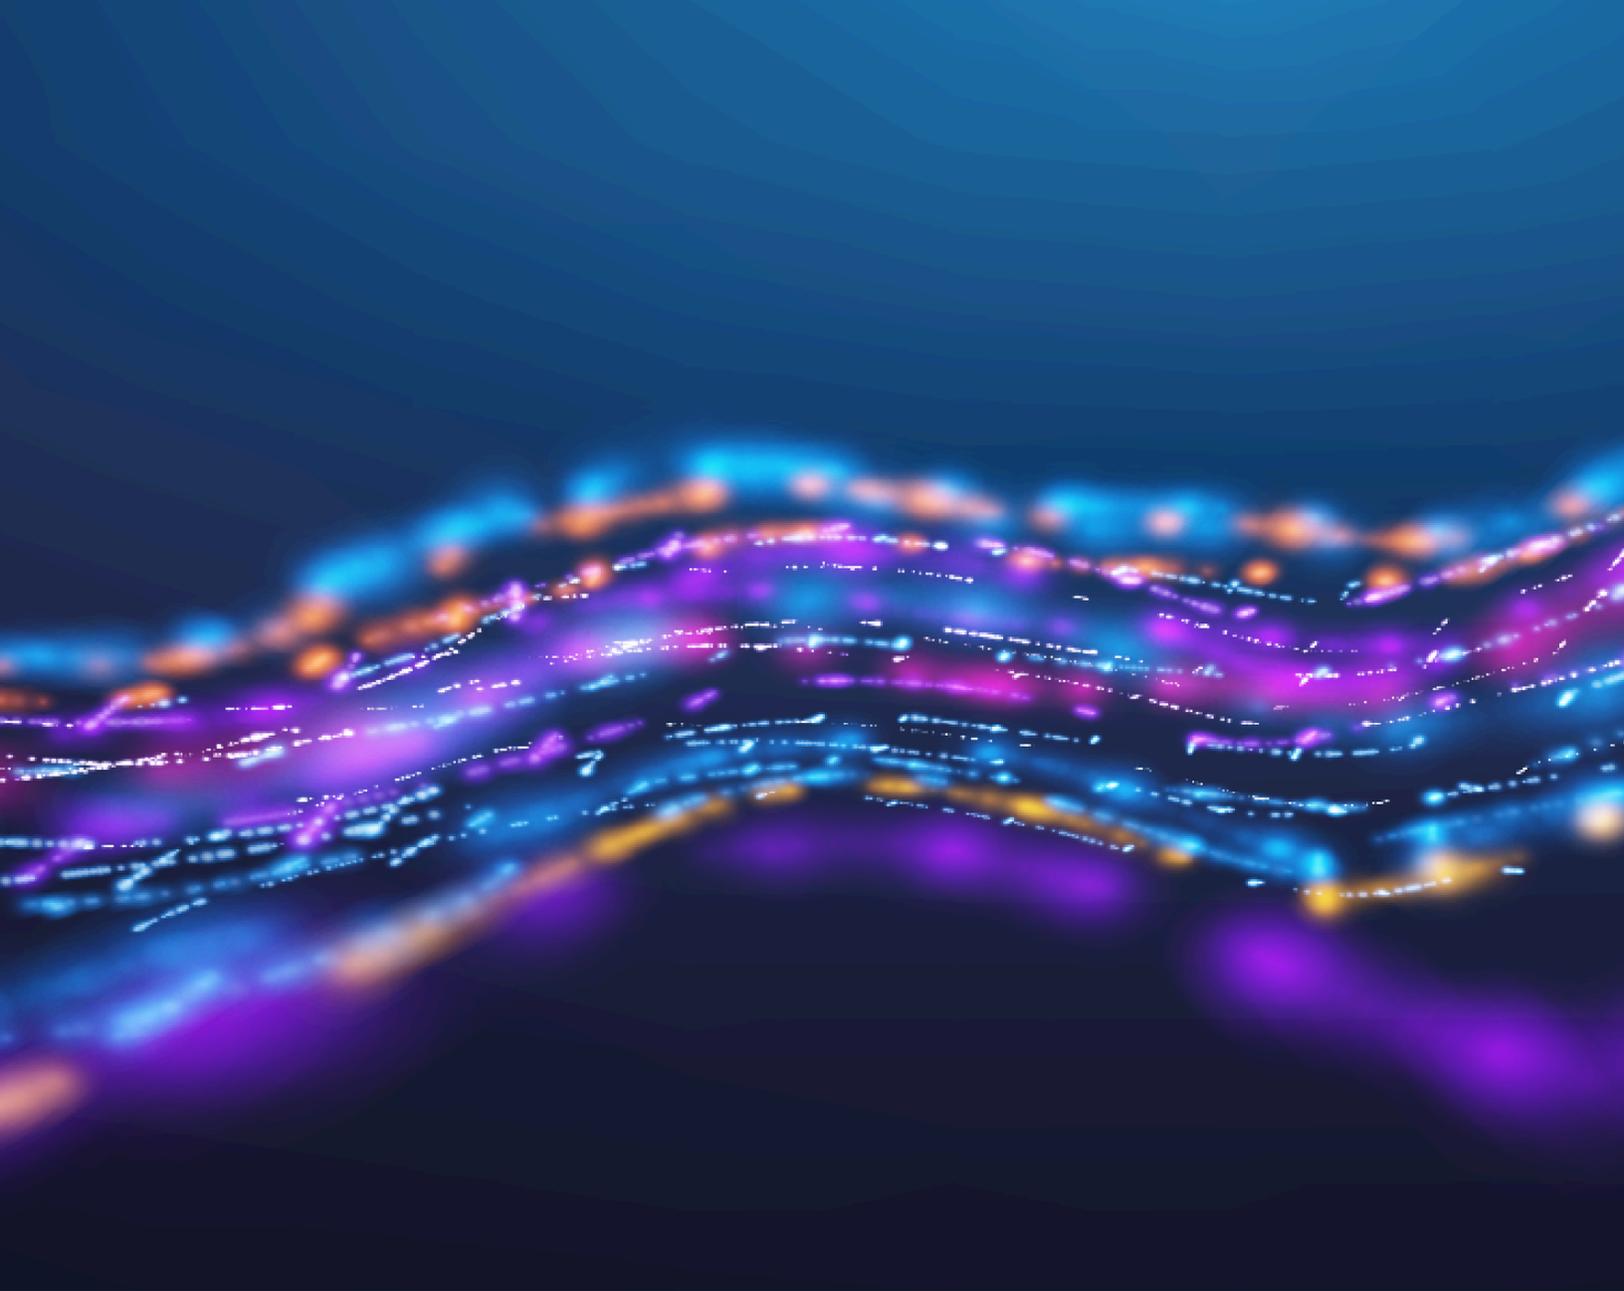
Adhere to the investigation framework

Organizations run their own investigations using the common framework, which is designed to leverage best practices from other sectors, enable a greater understanding of all the factors contributing to an incident, and create actionable findings — all contributing to continued safe innovation with GenAI.

Share insights with the community

While no organization is expected to share the details of incident investigations, lessons learned and insights uncovered from these investigations can contribute to the community's greater understanding of GenAI safety. Organizations are encouraged to share findings and insights so that everyone can continue to learn and innovate together.

If you'd like to learn more, see the full project, including incident definitions and examples and the GenAI Forward investigation framework, at aspenpolicyacademy.org/project/genai-forward-incidents-2026.



RDGraphics via Canva



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](https://www.aspenpolicyacademy.org).