

The MITRE logo is displayed in a bold, white, sans-serif font. It is positioned in the upper left corner of the slide, above a horizontal line that separates it from the tagline. The background of the slide features a stylized world map in orange and red, overlaid with a blue grid and various flight paths indicated by white and blue arrows and airplane icons.

**SOLVING PROBLEMS
FOR A SAFER WORLD®**

ADVANCING TOWARDS AUTOMATIC DETECTION OF SAFETY EVENTS

Generating safety intelligence to prevent aviation accidents

Stephen Szurgyi, The MITRE Corporation

December 2021

Safety assurance is not an accident

The Federal Aviation Administration (FAA) manages the world's largest, busiest, and most complex airspace in the world. More than 87,000 flights move through this airspace each day. These flights stay safe thanks to the continued collaborative efforts of FAA service providers and industry partners. The safety level of today's flight operations reflects a continual improvement in performance, technologies, and procedures. The goal of all aviation stakeholders is to make this system even safer.

Legacy approaches to safety provision have been compliance-based. This approach has been highly effective in getting us to the incredibly safe system we have today. However, the past decade has seen rapid advancements in aviation concepts and capabilities. The aviation system is in a discovery and rapid innovation phase for incorporating new entrants, controls, and business models. This evolving aviation system, as well as the role and responsibility of the FAA in ensuring continued safety, requires risk-based safety management. As the National Airspace System (NAS) gains increasing complexity via integrated operations, capabilities are needed to gather safety intelligence on systemic risks and the efficacy of safety barriers.

The generation of safety intelligence must occur continuously and proactively to provide maximum benefit. We must leverage innovative approaches that utilize artificial intelligence (AI) to identify safety issues and the underlying contributing factors (precursors) that lead to undesired aircraft states and system safety states. These AI-based technologies must support the discovery of safety issues through automatic identification of non-normal operational behavior.

SAFETY INTELLIGENCE IS

- **Actionable**
Fulfills defined need for insight into safety risk
- **Integrated**
Supports capabilities that are embedded in an established operational process
- **Reliable**
Trusted information supported by measured accuracy
- **Timely**
Delivered in time to enable mitigation of identified risk
- **Curated**
Vetted, customized source of verifiable information

There are many challenges to achieving this vision: safety culture, policy implications, human trust of AI output, and technological hurdles. To overcome many of these challenges, as noted by FAA Administrator Stephen Dickson, the FAA must be better at transitioning insights from its research partners into operational use¹. Now is the time to take action and convert advanced analytics research into the generation and operationalizing of safety intelligence.

¹ Talk given by FAA Administrator Stephen Dickson for the FAA Enterprise Information Management (EIM) Data Innovation Series, December 2020.

How to Operationalize Safety Intelligence

There are four key steps to enable effective generation of safety intelligence:

1. Collection and curation of appropriate aviation data
2. Detection of non-normal operational behaviors
3. Mapping the relationship between these non-normal behaviors and the impact they have on the safety of flight
4. Collaboration with stakeholders to validate those relationships, produce actionable (safety) intelligence, and develop mitigation strategies based on that intelligence.

(1) Data collection and curation

Data is foundational to detecting and learning from safety events and generating safety intelligence. Machine learning enhancements have enabled more effective utilization of historical data that can inform in-time monitoring of causal factors leading to safety events. Successful, continued operations of the NAS

“... our research partners [such as] MITRE, Volpe, NASA, and others, are doing some great safety analysis... if we get access to their data [and] dashboards, then we can move their research insights into operational use much faster.”

– Stephen Dickson, FAA Administrator

require interrelated management of shared data to improve aviation safety insights. Ubiquitous data can enable a much needed “learning environment” to sharpen anomaly detection.

Mapping data anomalies to safety issues requires access to multiple data repositories such as aircraft performance, weather, aviation risk data, aircraft position data, controller-pilot voice

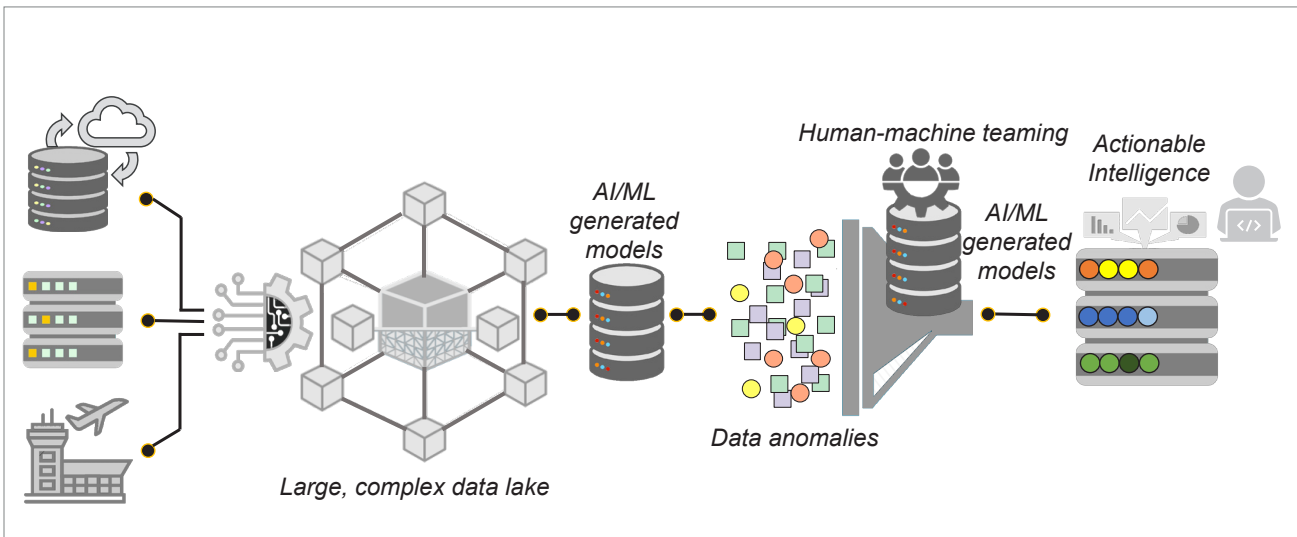


FIGURE 1. ILLUSTRATIVE PROCESS FOR TRANSLATING DATA ANOMALIES INTO SAFETY INTELLIGENCE

communications, human factors information, service outage data, and safety data reports. This data is often not available in an integrated and easily explorable manner.

Currently, many data sources across the FAA exist in silos. Differing processes, access restrictions, and priorities of the organizations that manage these silos limit the FAA's ability to capture the rich context surrounding safety events. As Administrator Dickson states, "FAA has a lot of data... [and] part of what contributes to stove-piping are these different data structures that we have... [and] one of the things that I want to do is look at some ways to bring our safety data and organizational data together in a more holistic way."²

FAA's Enterprise Information Management (EIM) system begins to overcome these organizational stovepipes. Transitioning critical data sets and analytic capabilities has begun, but EIM will take time to mature. FAA's Safety and Technical Training (AJI) Aviation Risk Identification and Assessment ([ARIA](#)) framework stands as an exemplar for this transition and, fueled by access to a broader range of operational, safety-related, and contextual data³, the Air Traffic Organization (ATO) and its research partners have begun to explore a range of complex safety topics: the effectiveness of ATO safety barriers; the impacts from degraded barriers for safety; the relationship between operational decisions and risk to NAS operations.

(2) Anomaly detection

Safety issues, and precursors to safety issues, are expected to be relatively rare and different from normal operations and, as such, show up as outliers in the data sets (i.e., as data anomalies).

“

FAA has a lot of data... [and] part of what contributes to stove-piping are these different data structures that we have... [and] one of the things that I want to do is look at some ways to bring our safety data and organizational data together in a more holistic way.

”

– Stephen Dickson, FAA Administrator

Anomaly detection is an active area of research. There currently exist multiple techniques, including those based on deep learning, to detect anomalies in data sets across various domains including aviation. Applying the current state-of-the-art approaches to large sets of aviation data can result in the detection of anomalies that are too numerous for manual vetting. As a result, anomaly detection is a key tool to advance aviation safety, as it can provide for proactive discovery of high-risk areas within the aviation system.

The ideal end-state is an AI-based solution detecting safety issues, identifying mitigations, and applying those mitigations with human supervision. Achieving this state will take time. Accelerating it will require significant investments in infrastructure to manage scalability and sustainability, subject matter expertise for training human-machine

² Fireside Chat with Peter F. Dumont and Stephen Dickson, Air Traffic Control Association (ATCA) Annual Conference and Exhibition, October 2019.

³ e.g., Mandatory Occurrence Reports (MORs), Electronic Occurrence Reports (EORs), historical Risk Analysis Events (RAEs), ARIA risk indices, Safety Key Performance Indicators (KPIs), traffic density and radar data, controller-pilot communications data, weather data.

teaming models, and process development to perform detection activity in a repeatable and sustainable manner.

FAA's EIM system can be an appropriate solution for instantiating appropriate AI infrastructure to enable effective anomaly detection through the application of machine learning models. Human-machine teaming, to develop those machine learning models, can train this anomaly detection capability to distinguish between data quality issues and those representative of deviations from normal operations. And, for this to be successful, mapping the relationships between anomaly profiles and safety issue precursors is crucial.

Until EIM can meet the scalability and sustainability requirements for continuous and comprehensive monitoring of the entire NAS, immediate impact can be achieved by targeting safety issue identification for individual areas of NAS operations. As these targeted applications are matured, best practices can be incorporated into an overall sustainable and repeatable process to monitor the entire NAS. This will also further foster the establishment of anomaly detection and mapping techniques, while providing for the development of processes and the growth of infrastructure necessary for continuous monitoring of the NAS.

(3) Mapping data anomalies to safety issues

Identification of safety issues from detected data anomalies presents some unique challenges related to resource constraints and the complexity of the domain. To overcome those challenges, a standardized and repeatable process for reviewing and mapping data anomalies to safety issues is needed. This includes:

- **Subject Matter Expertise.** Deep aviation subject matter expertise is necessary for timely review and validation of detected data anomalies

with safety implications. This primarily manual process is carried out by Subject Matter Experts (SMEs), which limits the speed at which review can occur.

- **Resource Burden.** Detected data anomalies are often too numerous for manual review by SMEs alone. Further, SME resources are also needed to develop and refine AI- and machine learning-based models to sift through complex sets of data anomalies.
- **Relationship to Safety.** Every data anomaly does not represent a safety issue. There may be underlying operating environment conditions that may result in a data anomaly that does not have any safety implications.

Through mapping these relationships, improvements to the underlying data sources can be identified and prioritized, and extended to new operating types and domains. By standardizing the process for mapping anomalies to safety issues, and applying machine learning techniques, SMEs' attention can more efficiently shift to assessing new and emerging anomalies, effectively mitigating the resource burden.

(4) Collaboration with stakeholders to produce actionable safety intelligence

The FAA's EIM system is intended to be the agency's centralized data lake of cross-organizational data and hub for integrated analysis capabilities. However, by making those data and advanced analytic capabilities accessible to a variety of experts across multiple stakeholder organizations, the FAA can better address complex safety issues through collaborative analysis. This collaboration concept, known as a collaborative research enclave (CRE), will provide the capacity to understand the interrelationship of operational anomalies surrounding high-risk events by synergizing multiple organizations' data pipelines.

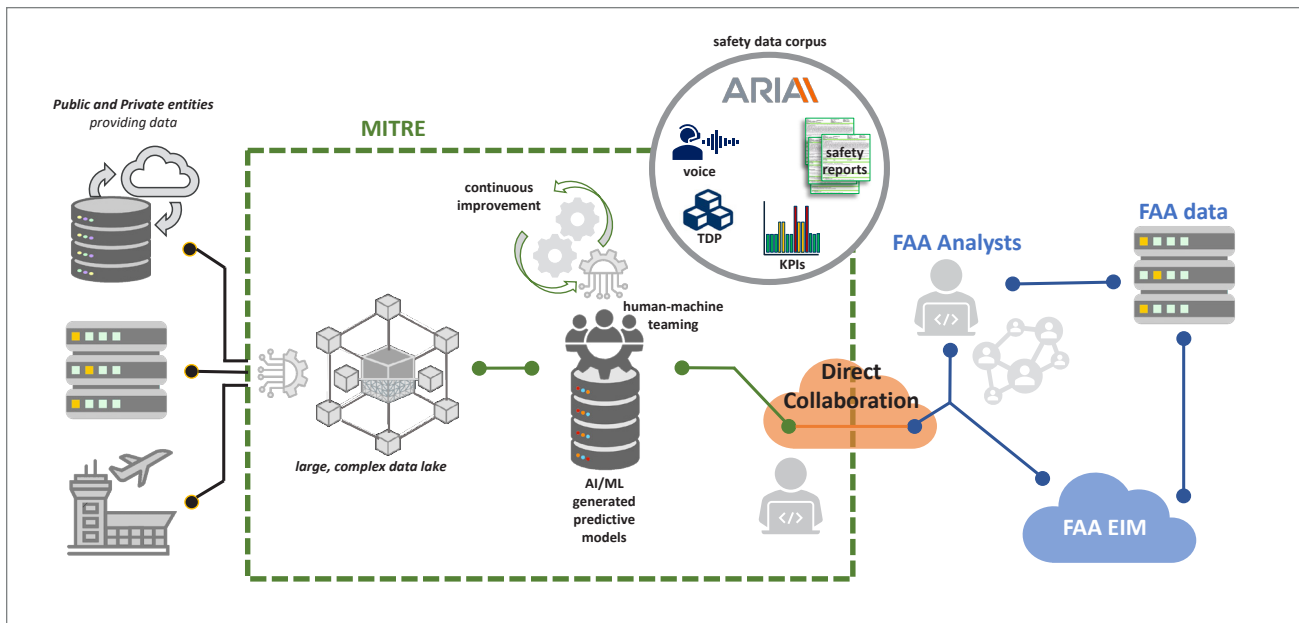


FIGURE 2. ILLUSTRATION OF COLLABORATIVE RESEARCH ENCLAVE CONCEPT

To reach the next level of safety, the agency should lean into connecting such environments to FAA's EIM, and implementing teams to mine for, identify, assess, and address systemic risks throughout the NAS. Given the complexity of these safety issues, these teams would best be comprised of diverse stakeholders responsible for the safety of NAS operations, including

representatives from airlines and airports, controllers, and pilots.

Leveraging safety intelligence-informed insights produced by these teams, the participating organizations internal and external to the FAA can take coordinated action to address the identified safety issues.

Recommendations

To make advancements in operationalizing safety intelligence by leveraging anomaly detection, aviation industry partners must:

- Increase the development of processes and tools that leverage human-machine teaming to evaluate integrated data for anomalies, which will require the elimination of current data stovepipes.
- Target specific opportunities for identifying the anomalies related to emerging risk(s) introduced into the NAS. This is necessary until continuous and comprehensive monitoring of the NAS is more viable. The introduction of new entrants (e.g., new airframes, unmanned aircraft system operations, and commercial space operations) offers a timely target as these new vehicles begin to integrate with and be adopted into routine NAS operations.
- Capture the mapping of data anomalies to associated safety issues, as well as the curation of that information into safety intelligence relevant to risk-based decision, in a standardized and widely available manner. We recommend the establishment of collaborative research enclaves to meet this need.

With gratitude for the support and contributions to this paper of Dr. Neeraj Koul.

About MITRE

MITRE's mission-driven teams are dedicated to solving problems for a safer world. Through our public-private partnerships and federally funded R&D centers, we work across government and in partnership with industry to tackle challenges to the safety, stability, and well-being of our nation.

The views, opinions, and/or findings contained herein are those of the author(s) and should not be construed as an official government position, policy, or decision unless designated by other documentation.



**SOLVING PROBLEMS
FOR A SAFER WORLD®**