

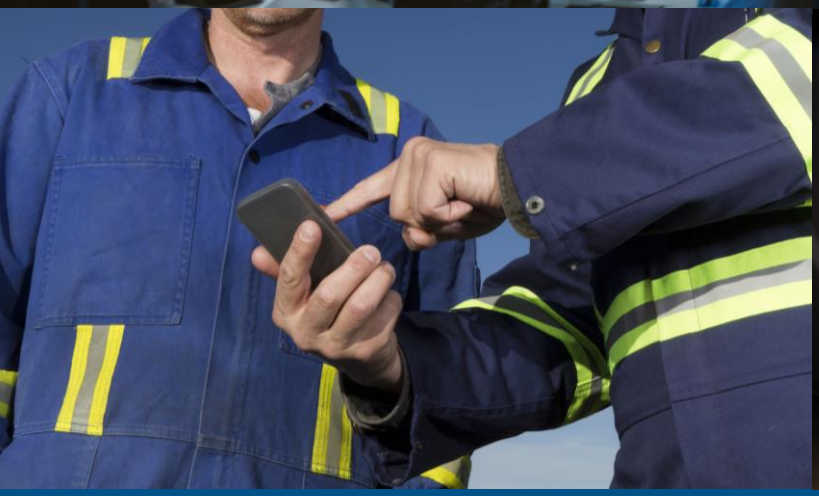


Emergency Management of Tomorrow Research: Summary Report

Leveraging AI for Emergency Management

December 2025

Science and Technology

This report was prepared for the U.S. Department of Homeland Security under a Work-for-Others Agreement with the U.S. Department of Energy, contract DE-AC05-76RL01830, IA 70RSAT23KPM000025

Pacific Northwest
NATIONAL LABORATORY

PNNL-38842

Emergency Management of Tomorrow Research: Summary Report

Leveraging AI for Emergency Management

December 2025

Jon Barr
Ann Lesperance
Rachel Bartholomew
Jessica Gray
Alex Hagen
Nick Betzsold
Kathryn Otte
Lissa Moore
Braden Sanders
Sara Smith
Ryan Poltermann
Maren Disney

U.S. DEPARTMENT of ENERGY

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

DISCLAIMER

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from
the Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062
www.osti.gov
ph: (865) 576-8401
fox: (865) 576-5728
email: reports@osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
or (703) 605-6000
email: info@ntis.gov
Online ordering: http://www.ntis.gov

PNNL-38842

Emergency Management of Tomorrow Research: Summary Report

Leveraging AI for Emergency Management

December 2025

Jon Barr
Ann Lesperance
Rachel Bartholomew
Jessica Gray
Alex Hagen
Nick Betzsold
Kathryn Otte
Elisabeth Moore
Braden Sanders
Sara Smith
Ryan Poltermann
Maren Disney

This report was prepared for the U.S. Department of Homeland Security under a Work for Others Agreement with the U.S. Department of Energy, Contract DEAC0576RL01830, IA 70RSAT23KPM000025.

Pacific Northwest National Laboratory
Richland, Washington 99354

About the Emergency Management of Tomorrow Research

The Department of Homeland Security (DHS) Science and Technology Directorate (S&T) is partnering with Pacific Northwest National Laboratory (PNNL) to execute the Emergency Management (EM) of Tomorrow Research (EMOTR) program to conduct research on strengthening and reimagining the future emergency response structure. The EMOTR program is identifying current EM research, eliciting capability needs from EM practitioners, and identifying where technology, such as artificial intelligence (AI), may benefit the future of EM and emergency operations centers. The project is delivering an iterative approach to curating a comprehensive framework to inform future research, development, and investments.

Executive Summary

A lack of effective data management approaches for situational awareness, interoperability, and integration to common operating pictures hampers EM by limiting effective communication, data sharing, resource allocation, and real-time decision-making. Currently, the EM community needs to be able to leverage the myriad EM data sources and data types to enable systems that can utilize real-time data, geospatial information, and AI tools to support EM for local and state governments.

To address this enduring challenge, the EMTOR program, led by DHS S&T and PNNL, engaged EM and public safety stakeholders and the private sector to identify current approaches that support enhanced situational awareness and common operating environments with AI enablement. This effort employed a structured methodology to characterize EM environments, identify data-related processes and interactions, and assess gaps, challenges, and opportunities surrounding AI integration to inform actionable insights that address current and future needs in EM operations.

As a result, this report includes a series of technology profiles that aim to mitigate challenges surrounding the use of EM data by AI systems. Featured technologies and approaches highlighted in this report include the following:

- Infrastructure
  - Data storage architectures
  - Data access architectures
- Agent Internet
  - Autonomous systems
  - Agent schema standards
- Protocol
  - Agent-to-agent protocol
  - Model context protocol
  - Agent gateways
  - Data interoperability standards
  - Data provenance and trust frameworks
- Tooling and Enrichment
  - Data discovery for improved accuracy: Fine tuning, retrieval-augmented generation
  - Data quality and safety controls for AI systems
  - Privacy-preserving collaboration and data sharing
  - Real-time intelligence and continuous model enhancement
- Memory and Personalization
  - Handling missing and/or degraded data quality
  - Handling missing data modalities
  - Controlling protocols for sudden retroactive access to data
  - Multi-jurisdictional knowledge access and compliance
  - Contextual knowledge integration for enhanced decision support
- Governance and Operations
  - Real-time data ingestion services
  - Data sovereignty and cross-jurisdictional compliance
  - Accountability, audit, and incident response for AI systems
  - Trust and quality assurance for AI operations

Each profile provides an overview of the concept, considerations, and opportunities for each of the following stakeholders: the EM community, policymakers, EM technology vendors, and AI leaders.

Given the complexities and rapidly moving state of disasters, applying AI to the EM discipline is an area primed for innovation. This report is designed as an informative tool for EM stakeholders seeking to leverage the power of large language models and AI reasoning systems to effectively plan for, respond to, and recover from crises and disasters.

Acknowledgments

PNNL would like to thank the following individuals and their numerous communities and partners for providing their perspective and input to further benefit the EMOTR project, including:

- Emily Martuscello, City of Nashua and Rivier University
- Mike Chard, Boulder County and the Boulder County Sheriff's Office
- Curry Mayer, Seattle Office of Emergency Management and the Seattle Emergency Operations Center
- Carrie Speranza, International Association of Emergency Management
- Grant Tietje, Front Range LLC
- Kamran Atri, DHS contractor
- Michael Simcock, DHS contractor
- Lynn Martin, Chris Sadler, Virginia Innovation Partnership Corporation
- Scott Drew, ATA Aviation
- Numerous additional EM and public safety officials supported workshops, exercises, and conferences throughout the EMOTR program.

Portions of this document utilized the assistance of PNNL's secure generative AI platform. The final content was reviewed, edited, and approved by the PNNL EMOTR team for accuracy and relevance.

Acronyms and Abbreviations

A2A	Agent to Agent
AI	Artificial Intelligence
C2PA	Coalition for Content Provenance and Authenticity
CAD	Computer-Aided Dispatch
COP	Common Operating Picture
DevOps	Development Operations
DHS	Department of Homeland Security
DHS S&T	Department of Homeland Security Science and Technology Directorate
DLP	Data Loss Prevention
DMSR	Data Management, Sharing, and Retrieval
EI	Edge Intelligence
EM	Emergency Management
EMOTR	Emergency Management of Tomorrow Research
EOC	Emergency Operations Center
FinOps	Financial Operations
GIS	Geographic Information Systems
HAZMAT	Hazardous Material
HRL	High Human Readiness Levels
ICS	Incident Command System
LLM	Large Language Model
MCP	Model Context Protocols
MOU	Memorandum of Understanding
ML	Machine Learning
NIEM	National Information Exchange Model
NIMS	National Incident Management System
PII	Personally Identifiable Information
PNNL	Pacific Northwest National Laboratory
RAG	Retrieval Augmented Generation
RMS	Resource Management System
SAFECOM	Public Safety Communications
SBOM	Software Bill of Materials
SDK	Software Development Kit
SPDX	Software Package Data Exchange
TRL	Technological Readiness Levels

Contents

Figures

1.0 Introduction

The Department of Homeland Security (DHS) Science and Technology Directorate (S&T) partnered with Pacific Northwest National Laboratory (PNNL) to execute the Emergency Management (EM) of Tomorrow Research (EMOTR) program to strengthen and reimagine the future of emergency response. EMOTR is a research and development effort focused on increasing the capacities of state and local public safety officials to plan for, respond to, and recover from emergencies of all types and sizes. To do this, the program leverages innovative approaches to engage EM and public safety stakeholders, the private sector, and interagency collaboration to enable rapid research to operations. This document shares a suite of opportunities and considerations for EM stakeholders, including emergency managers, policymakers, EM vendors, and artificial intelligence (AI) leaders, to advance AI-enabled technologies for EM. Future opportunities are highlighted in Section 6.0 and the technology profiles located in Appendix A.

1.1 Problem

A lack of effective data management approaches for situational awareness, interoperability, and integration to common operating pictures (COPs) hampers EM by limiting effective communication, data sharing, resource allocation, and real-time decision-making. Currently, the EM community needs to be able to leverage the myriad EM data sources and data types to enable systems that can utilize the breadth of real-time data, geospatial information, and AI tools to support EM for local and state governments.

This report investigates technologies, techniques, and architectures to effectively utilize multimodal data for integration into large language models (LLMs) for homeland security. As a part of this research, this report identifies current tools and challenges to enhance data interoperability through standards, algorithms, and architectures.

1.2 Background

This effort builds on previous EMOTR activities that assessed the EM and AI landscapes, elicited stakeholder feedback, and ultimately drove findings into recommendations for how technology, such as AI, may benefit the future of EM and Emergency Operations Centers (EOCs). Most recently, in 2025, three exercises convened EM and public safety practitioners to brainstorm the next generation of technology solutions needed to address high-priority challenges encountered in EOCs and public safety. The exercises gave specific attention to assessing the technologies' impact and feasibility. Results and the impact of these exercises are captured in a separate report, “Emergency Management of Tomorrow Research: EOC of the Future Technology Exercises” (September 2025).[1]

A key issue that surfaced from the exercises, as well as in prior engagements with the EM community[2] from across the country, is that the lack of effective situational awareness among all stakeholders hinders EM response and recovery operations. This is an issue because individual systems and processes leave significant gaps in consistent and timely communication, data and

[1] Betzsold N.J., A.K. Otte, J.L. Barr, A.M. Lesperance, R.A. Bartholomew, J.A. Gray, and C.M. Sleiman, et al. 2025. *Emergency Management of Tomorrow Research: Emergency Operations Center of the Future Technology Exercises*. PNNL-38245. Richland, WA: Pacific Northwest National Laboratory.

[2] For purposes of this report, our primary focus is the EM community, but in some cases, this is applicable to the broader public safety community.

information sharing, resource allocation, and real-time decision-making. EM stakeholders expressed a need for a solution that incorporates real-time data feeds, geospatial information, and AI, available to all local and state governments. This would effectively be a reimagining of the foundation for a COP that could create a national network for intelligent data sharing, dramatically improving the effectiveness of response and recovery operations from the local to federal level.

AI, machine learning (ML), and agentic solutions have the potential to revolutionize situational awareness by integrating and harmonizing data across disparate systems, stakeholders, and geographies. Leveraging advanced AI agents could deliver a cost-effective, scalable, and accessible framework that equips emergency managers and public safety officials with real-time insights and a unified COP, which aligns with the Office of Management and Budget and Office of Science and Technology Policy Memo FY 2027 Administration Research and Development Priorities and Cross-Cutting Actions that states: "*agencies should explore novel datasets and identify opportunities to leverage AI applications to enable improved resilience*".[1]

1.3 Objectives

To address these enduring challenges and opportunities for improvement, PNNL conducted research to identify current approaches that support enhanced situational awareness and common operating environments with AI enablement. This was accomplished through the following efforts:

- Investigate technologies, techniques, and architectures to effectively utilize multimodal data for integration into LLMs for homeland security.
- Identify current tools and opportunities to enhance data interoperability through standards, algorithms, and architectures.

Together, these efforts employed a structured methodology to characterize EM environments, identify data-related processes and interactions, and ultimately assess gaps, challenges, and opportunities surrounding AI integration to inform actionable insights that address current and future needs in EM operations.

1.4 About this Report: Value Proposition and Use

This document supports various stakeholders of the greater EM community, from emergency managers to the private sector. This report details numerous opportunities for the private sector to play a critical role in innovating and deploying tools into the EM discipline—an area that is primed for innovation—all with the aim of enabling a shared public-private sector vision for seamless information sharing.

[1] Executive Office of the President. 2025. *Fiscal Year (FY) 2027 Administration Research and Development Budget Priorities and Cross-Cutting Actions (M-25-34, NSTM-2).* Washington, DC: Office of Management and Budget. Accessed October 2023. https://www.whitehouse.gov/wp-content/uploads/2025/09/M-25-34-NSTM-2-Fiscal-Year-FY-2027-Administration-Research-and-Development-Budget-Priorities-and-Cross-Cutting-Actions.pdf

1.4.1 Audience and End Use

This report is designed for the EM community, policymakers, EM technology vendors, and AI leaders seeking to leverage the power of LLM and AI reasoning systems to enhance state and local officials' ability to effectively plan for, respond to, and recover from crises and disasters.

This report includes a series of user-friendly technology profiles that aim to mitigate challenges surrounding the use of EM data by AI systems (see Appendix A). Each profile provides an overview of the concept followed by considerations and opportunities for the following stakeholders: the EM community, policymakers, EM technology vendors, and AI leaders.

- For the *EM community*, the profile details what the specific concept will enable them to do, what they can currently do relative to the concept, and what they need to do to prepare for incorporation into their systems or COP.
- For the *policymakers*, the profile provides a policy-focused overview and opportunities to advance with regard to usage, governance, and other key policy decisions required for proper policy-aligned usage.
- For *EM vendors*, the profile provides a more technical overview of the concept and its usage, and opportunities for vendors to incorporate or take advantage of the concept in their products. This information provides specific opportunities for the private sector and developers to innovate.
- For *AI leaders*, the profile details the nuances of applying these concepts to the EM domain and opportunities to adopt these concepts in this domain.

1.4.2 Understanding AI-Enabling Technology: Focus and Intent

This document summarizes how AI-enabling technologies fit together and work as a system. Why is this important?

By demonstrating the interconnected nature of AI-enabling technologies, organizations ensure they leverage the full potential of AI, create systems that are transparent and efficient, and foster collaboration between all stakeholders. It is important to demonstrate how AI-enabling technologies fit together and work as a system for several reasons, including:

- Holistic Understanding – AI systems typically involve multiple interconnected components, such as machine learning models, cloud computing, data storage, and APIs. Demonstrating how these technologies collaborate helps stakeholders understand the bigger picture, including how each technology plays a critical role in achieving desired outcomes. This fosters clarity and trust in AI solutions.
- Optimized Integration – AI technologies do not function in isolation; they rely on integration between tools and platforms. Showing how they fit together highlights dependencies, interoperability, and potential integration challenges. It ensures proper system architecture design and efficient workflows, leading to better performance and reliability.

- Scalability and Sustainability – Understanding the interconnectedness of AI-enabling technologies enables developers and organizations to design scalable and sustainable systems. For example, integrating data pipelines, AI models, and deployment infrastructure as a unified system allows future upgrades and expansions to be more efficient.
- Improved Collaboration Across Teams – Different teams, including data scientists, engineers, business analysts, and IT specialists, often contribute to building AI systems. Demonstrating how AI technologies work together aligns these teams with shared objectives and ensures smooth collaboration by clarifying roles and responsibilities within the system.
- Transparency and Trust – AI systems can seem abstract or complex to end users. Explaining how the components interact helps demystify the technology and builds trust, particularly in industries that require ethical considerations, such as healthcare or finance. Transparency in AI systems also helps mitigate biases and risks.
- Streamlined Troubleshooting and Problem Solving – When AI technologies are understood as a cohesive system, it becomes easier to identify where issues or inefficiencies arise.
- Facilitates Knowledge Sharing – When AI systems are documented and showcased as systems of interconnected technologies, they promote knowledge sharing among developers, researchers, and industry practitioners. It enables faster innovation and adoption of best practices across ecosystems.
- Navigating Complexity – AI systems can be incredibly complex, with components like neural networks, natural language processing, data pipelines, inference engines, and monitoring tools working together. Showing how these fit together breaks down the complexity and makes it easier for teams and organizations to adopt, modify, or enhance the system.

When talking about AI-enabling technologies, what types of technologies do you mean?
To accomplish the objectives detailed in Section 1.3, this effort examines the following technologies and approaches that enable AI applications. This includes the following categories of the AI technology stack (detailed in Section 2.0):

- Infrastructure (specifically core platform services [e.g., data storage approaches])
- Agent Internet
- Protocols. (e.g., example Agent to Agent [A2A] protocols, Model Context Protocols [MCP])
- Tooling and Enrichment
- Memory and Personalization
- Operations and Governance

Additionally, some key non-technical considerations around operations and governance are addressed in Section 4.0, Foundational Considerations for AI Adoption in EM.

A major part of this document is the Appendix, which details the technology profiles. What are they and why are they important?
The technology profiles (available in Appendix A) create a baseline of information about the technology product and can be used as a reference. Technology profiles capture key information on a given domain (e.g., data fabrics, MCPs) and identify specific technologies and articulate how they “fit “into a larger AI system. These profiles also include a narrative about specific technology and how various EM stakeholders can use it. These stakeholders include *EM community*, *policymakers, EM vendors,* and *AI leaders*.

2.0 Methodology

The dynamic and multifaceted nature of EM necessitates a robust approach to integrate innovative technologies, such as AI, to support enhanced planning, response, and recovery operations. To address this need, the EMOTR team developed and implemented a structured methodology to align AI solutions with the operational realities of EM. This methodology has five interconnected activities that characterize the EM environments, identify data-related processes and interactions, select an appropriate AI-aligned framework, analyze relevant technologies, and assess gaps, challenges, and opportunities surrounding AI integration (Figure 1). By leveraging expert insights from EM subject matter experts and AI/ML professionals, this work provides a comprehensive foundation for integrating AI technologies effectively into EM contexts while ensuring safety, security, and efficiency. The following sections detail the EMOTR methodology, providing clarity on how scenarios, frameworks, and analyses were used to inform actionable insights that address current and future needs in EM operations.

2.1 Characterize Emergency Management Operational Environments

To ensure this work aligned with EM needs, it was imperative to capture the complexities of the operational environment in order to validate technologies around real-world events. Two scenarios were developed to bring out those complexities:

- Damage assessment and service restoration after an earthquake, and
- Terrorist attack involving a swarm of autonomous unmanned aerial systems (UAS).

Each scenario was examined from two perspectives: 1) from a well-staffed and resourced urban EM department and operations center, and 2) from a small rural EM department and operations center with limited staffing and resources. These scenarios provided an opportunity to identify numerous data types, owners, pathways, and needs that could be supported by current and future AI systems.

2.2 Identify Processes and Interfaces Associated with Data Management of AI Use

Working with a team of EM subject matter experts along with AI/ML experts with backgrounds in systems engineering, data science, and computer science, the scenarios detailed in Section 2.1 were analyzed to determine the various processes/activities and interfaces that AI technology could interact with, utilize, or handle (interaction points) to provide EM more effective planning, response, and recovery. These data interaction points were then assessed for what would be required for safe, secure, efficient, and effective data management of these AI systems and services. This set of data, interaction points, and needs provided the foundation for surveying the AI technology landscape for EM solutions that could be applied.

2.3 Select Emergency Management-Aligned Framework

Selecting an EM-aligned AI framework for this effort allowed for a common framing and language to communicate between the various stakeholders (EM community, policymakers, EM vendors, and AI leaders). Additionally, the framework provided longevity to this work as AI technology and techniques are advancing at an incredible pace. Such a framework provides a functional structure that will likely be maintained even as the technologies associated with the different functions advance and change. This will also allow stakeholders using this report to

both understand the technologies that can be leveraged at the time of the writing and find technologies associated with the functional elements of the framework going forward.

2.4 Investigate Technologies, Techniques, Services, and Architectures

The EMOTR team examined current AI technologies, techniques, services, and architectures aligned with the framework discussed above and detailed in Section 3.0. The AI technologies, techniques, services, and architectures were analyzed and then structured into technology profiles (Appendix A) organized by framework functional domains. Each technology profile contains an overview followed by stakeholder-specific analysis. The overview is intended to be informative and accessible by all stakeholders. The stakeholder-specific analysis ensures:

- *EM and other operational stakeholders (emergency managers)* know why this is relevant to their work and how it can improve operations. This analysis is informed by the application of technology to scenarios developed in Section 2.1.
- *Standards and policymakers* understand the current state of governance, standards, and best practices are associated with implementing these solutions across the nation.
- *EM vendors* have a more technical understanding of what the solution does and how it can be implemented.
- *AI leaders* understand the complexities, requirements, and constraints that will be relevant when accessing their models and systems for EM activities.

The EMOTR team also conducted a limited market landscape assessment on EM solutions focused on integrating data lakes, data lake houses, data mesh, and data fabric architectural approaches to gather insight into market trends and vendor capabilities.

2.5 Assess Gaps, Challenges, and Opportunities

This effort also examined the gaps, challenges, and opportunities associated with integrating these technologies in the EM space. The following sections detail various aspects of this analysis as follows:

- Section 4.0 details a series of core considerations that will help ensure EM organizations are prepared to effectively use AI in their work.
- Section 5.0 and Appendix A highlight promising technologies and approaches, as well as specific gaps, challenges, and opportunities associated with each technology.
- Section 6.0 details opportunities for EM stakeholders to advance and utilize AI systems for EM operations.

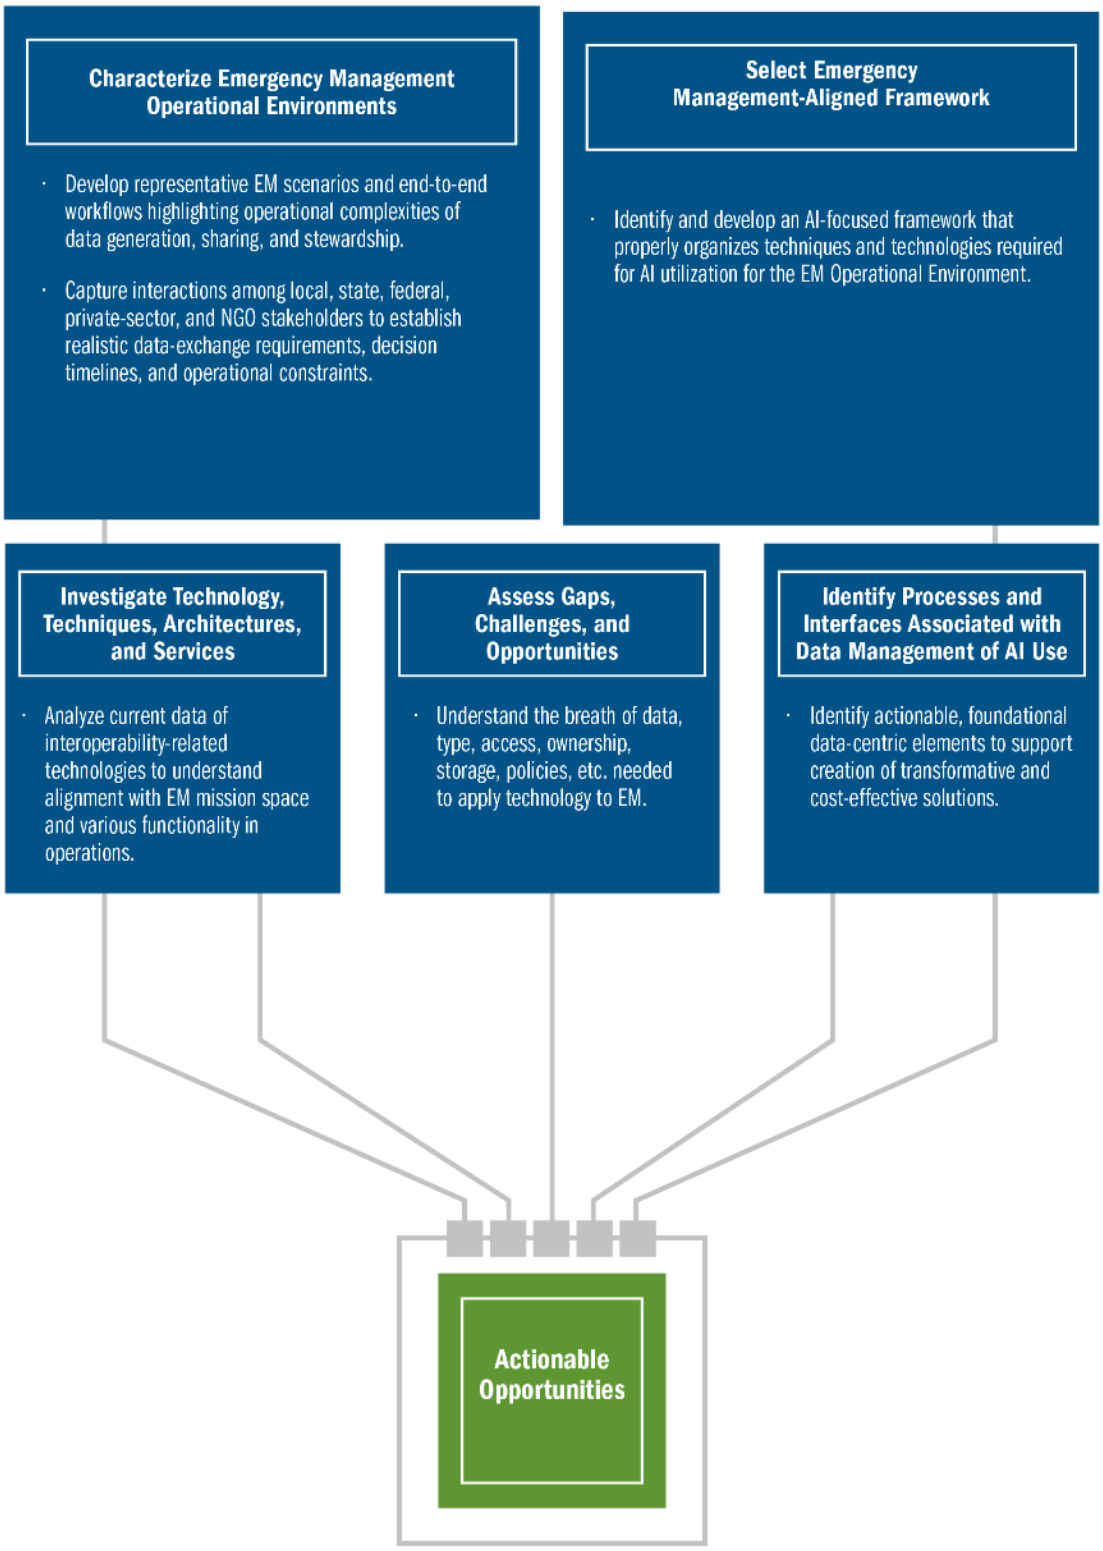

Figure 1. Methodology to identify key considerations to leverage AI systems and tools for EM.

3.0 Organizing Structure and Framework

This report organizes the technologies and capabilities needed to meet the data, operational, and security requirements of the EM community, enabling full advantage of both commercial and open-source AI systems (including agentic systems and LLMs). Profiles are grouped by layers of the AI stack (see Figure 2). Although the focus of this activity was on enabling AI-driven workflows, this document does not prescribe specific AI behaviors. Instead, it concentrates on the end-to-end considerations required to discover relevant data, securely access and move data, and integrate data with appropriate AI-enabled systems.

Subsequent sections of this report describe each layer, why the layer is important, how it relates to other layers, and a series of profiles of technologies and capabilities that fall under that functional umbrella. These profiles provide the current foundational elements that can be utilized to enhance AI-enabled tools and break down roadblocks to safe and effective information identification, retrieval, access, and utilization.

As discussed previously, layers of the framework contain relevant technologies that directly support a holistic AI capability for EM. Figure 3 details how these technologies could currently be structured into a system that supports EM. This figure examines potential general AI functionality applied to one aspect of the earthquake validation scenario discussed in Section 2.1, specifically ensuring water is delivered to where it is needed. Along the top of the figure are eight steps that the AI system advances through in order to support the initial request by EM and provide a proper response. These steps highlight the supporting pieces of the framework included in the conceptual architecture, providing an applied and grounded view of the framework.

Every technology profile in a framework layer includes a conceptual overview and stakeholder-focused content. The concept overview includes an explanation of the technology or capability, and, as appropriate, critical considerations including potential challenges, limitations, and best practices, and stakeholder-tailored opportunities for action.

By following this structure, readers can quickly identify the technologies relevant to their role, understand how those technologies interconnect, and take steps toward secure, effective, AI-enabled EM operations.

Layers of AI Stack	Description
1. Infrastructure	Provides the foundational compute, storage, and network resources (on-premises, cloud, or edge) that power every other layer. Its purpose is to deliver resilient, scalable, and cost-optimized hardware and core platform services without prescribing any particular vendor or technology.
2. Agent Internet	Establishes the universal connectivity fabric through which agents discover, access, and interact with external data sources, services, and one another. It abstracts transport, authentication, and service-discovery concerns so agents can operate across heterogeneous environments.
3. Protocols	Defines the common languages and interaction patterns that allow agents to communicate, negotiate tasks, share context, and enforce trust. By standardizing message formats, capability declarations, and security primitives, it ensures interoperability while remaining agnostic to specific implementation stacks.
4. Tooling and Enrichment	Supplies modular extensions (e.g., function call wrappers, data transformers, and execution sandboxes) that expand or refine an agent's native abilities. This layer enables rapid incorporation of new skills and resources without altering core reasoning components.
5. Cognition and Reasoning	Implements the planning, deliberation, and decision-making logic that turns goals into executable action plans. It encompasses task decomposition, feedback loops, and adaptive strategies while remaining independent of any one algorithmic approach.
6. Memory and Personalization	Maintains short- and long-term contextual knowledge, user preferences, and experiential learning that allow agents to adapt over time. By separating storage and retrieval mechanisms from specific databases or models, it supports future evolutions in memory technology.
7. Application	Hosts end-user and domain-specific solutions, ranging from conversational assistants to complex multi-agent workflows, that leverage the underlying cognitive and tooling layers. It focuses on orchestrating capabilities into coherent experiences without binding them to a particular user interface or platform.
8. Operations and Governance	Oversees deployment, monitoring, compliance, security, and policy enforcement across the entire stack. It provides guardrails, auditing, and life-cycle management to ensure responsible, reliable, and cost-effective operation independent of specific DevOps or governance tooling.

Figure 2. Selected AI stack and organizational framework.

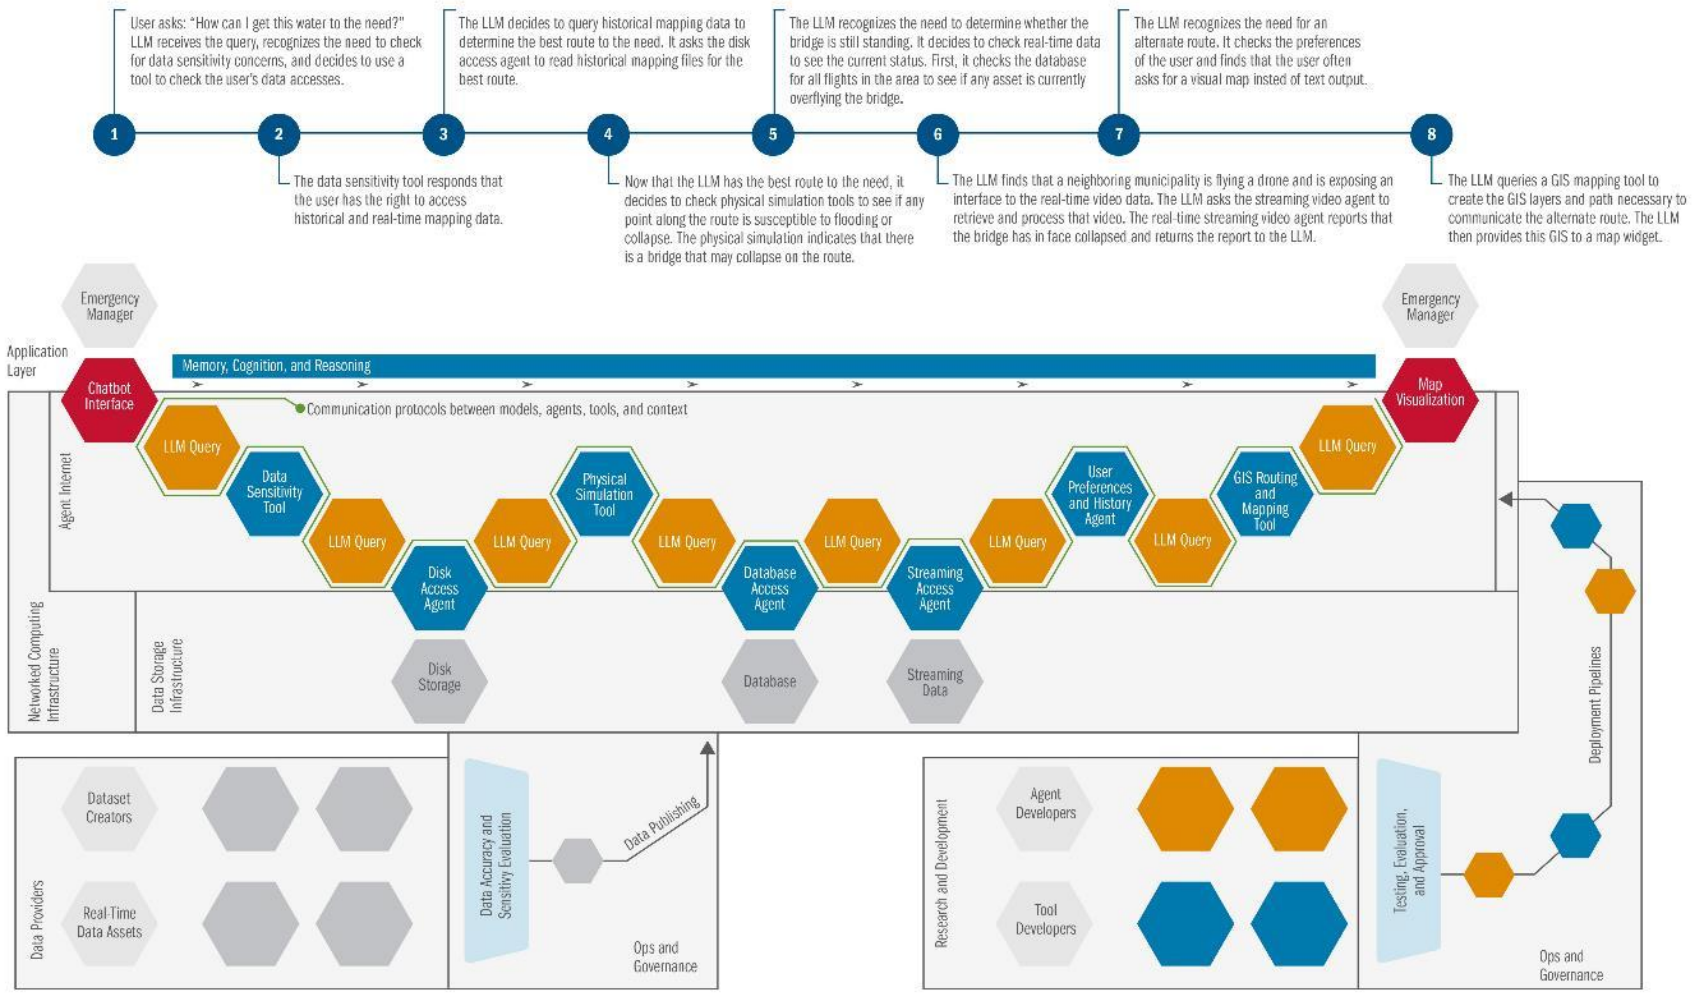

Figure 3. Detail of a conceptual EM architecture leveraging the AI stack.

4.0 Foundational Considerations for AI Adoption in EM

The following sections detail potential barriers that, unless resolved, will prevent any organization from taking full advantage of generative AI's potential and may even lead the organization to experience widespread, long-term, negative impacts. Preparation is needed to align an organization's people and processes with effective, efficient, and ethical generative AI usage without causing unintended consequences. Beyond the foundational considerations for AI adoption detailed in this section, Appendix B provides lessons learned and best practices associated with developing, implementing, and operating these AI systems in national security settings.

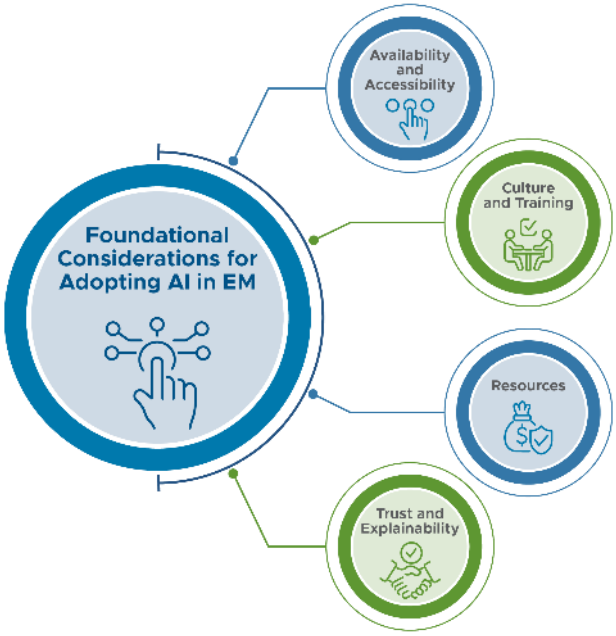

4.1 Availability and Accessibility

Availability refers to whether the necessary high-quality data is available digitally, while accessibility concerns how easily and quickly these data can be retrieved, processed, and presented to an AI to inform task completion. High-quality data must be provided to AI tools for training, tuning, and response creation efforts in a timely manner.

Potential Barrier: Are the necessary training data available digitally?	
Why It Matters	**What EM Should Know**
• LLM hallucinations[1] are a barrier to their usage in high-consequence environments. • High-quality training data can mitigate hallucination in LLM responses. • LLM practitioners identified 13 attributes of high-quality LLM training data: reliability, relevance, accuracy, compliance, accessibility, privacy protection, documentation, large-scale data, diversity, knowledge content, wide range of sources, absence of low-quality documents, and absence of toxic data.[2] • Edge intelligence (EI) can bolster dataset comprehensiveness while mitigating biases	• Gaps in disaster-related datasets skew model knowledge about those events.[4] • Data on disaggregated local records of impacts of smaller-scale events can help better understand incident consequences.[4] • Model training data must overcome existing "disruption to communication channels, infrastructure, and the health system [that] create barriers to information sharing and governance" if EM practitioners want to understand and predict incident economic losses and psychosocial damage.[4] • Ideal EM training datasets would be long-term, comprehensive, externally validated, and standardized.[4]

[1] Hallucinations occur when a model's response contains inaccurate, made-up, or nonsense information.

[2] Yu, X., Zhang, Z., Niu, F., Hu, X., Xia, X., & Grundy, J. (2024). What makes a high-quality training dataset for large language models: A practitioners' perspective. *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*, 656–668. https://doi.org/10.1145/3691620.3695061

[4] Guha-Sapir, D., Scales, S. E., Kruczkiewicz, A., Arfvidson Umaña, N. M., Das Gupta, I., Colaço, C., & Venkatraman, R. (2023). (tech.). *Closing Climate and Disaster Data Gaps: New challenges, new thinking*. Geneva, Switzerland: UNDRR.

Potential Barrier: Are the necessary training data available digitally?	
Why It Matters	**What EM Should Know**
and underrepresentation by aggregating data from local devices.[1] • Introducing mathematical tools into LLM data pipelines helps identify where datasets need quality augmentation.[1]	• Future prediction accuracy is aided by training data recorded as close as possible to the time and place of previous incidents.[4]

Potential Barrier: Are the necessary use-case (test) data available digitally?	
Why It Matters	**What EM Should Know**
• LLMs cannot access analog data. • Data must be digitized for LLMs to access and incorporate into tasks and responses. • Many historical sources remain undigitized.[2] • Content currently in forms such as paper notes and conversations is inaccessible to LLMs without intermediate conversion tools.[2]	• EM data often exists as paper notes, sticky notes, or conversations between practitioners.[3] • LLMs will require intermediate tools for digitization and transcription efforts to convert these data into a usable format.[3] • EM field methods and tools currently limit digitization during incidents.[3] • Mobile phones and tablets can enable quick generation and sharing of digitized information on the go and in the field.[3] • Digitization also improves data transfer timeliness, enhancing situational awareness.[3]

Potential Barrier: Is the required use-case data accessible within an appropriate timeline?	
Why It Matters	**What EM Should Know**
• Traditional predictive models depend on massive historical datasets and may miss identifying emerging patterns.[4] • LLMs with few-shot or zero-shot learning capabilities can outperform traditional models on certain predictive tasks.	• Timeliness is crucial in EM as incident conditions change rapidly. • Samsara survey of EM professionals across 21 industries in 7 countries, including the U.S., demonstrates:[5]

[1] Zhang, X., Xie, G., Huang, Y., Xiong, Z., Liu, J., Cui, S., Sun, S., & Sherman Shen, X. (2025). Edge intelligence in the Generative Artificial Intelligence Era. *IEEE Wireless Communications*, 1–9. https://doi.org/10.1109/mwc.2025.3599652

[2] Baack, S., Biderman, S., Odrozek, K., Skowron, A., Bdeir, A., Bommarito, J., Ding, J., Gahntz, M., Keller, P., Langlais, P.-C., Lindahl, G., Majstorovic, S., Wolf, T., Kydlíček, H., Gutermuth, L., Fadaee, M., Chmielinski, K., Belião, J., Baker, M., … Leppert, G. (2025). Towards Best Practices for Open Datasets for LLM Training. *Proceedings from the Dataset Convening*. https://doi.org/10.48550/arXiv.2501.08365

[3] Panasonic Connect Blog. (2025, September). *Why it's time for fire EMS to ditch pen and paper*. Panasonic North America - United States. https://connect.na.panasonic.com/blog/toughbook/why-its-time-for-fire-ems-to-ditch-pen-and-paper

[4] Song, Y., Ma, R., Liu, J., Guo, Y., & An, S. (2025). How can timeliness and accuracy be balanced? the surveillance information reporting mechanism for public health emergencies based on the Online Direct Reporting System. *BMC Public Health*, *25*(1). https://doi.org/10.1186/s12889-025-24637-5

[5] Government Fleet (2025, October). *New Samsara data reveals dangerous gap in disaster readiness*. Government Fleet – United States. https://www.government-fleet.com/10248923/new-samsara-data-reveals-dangerous-gap-in-disaster-readiness

Potential Barrier: Is the required use-case data accessible within an appropriate timeline?	
Why It Matters	**What EM Should Know**
• LLM-based prediction requires timely, relevant use-case inputs. • Outdated data lowers LLM accuracy and output quality in dynamic situations.	○ 95% of organizations face financial losses during crises due to the inability to locate assets, ○ 64% of organizations lack consistent real-time access to key operational data, and ○ 79% of organizations fear losing communications if infrastructure is damaged. • EM professionals report frustration over real-time intelligence gaps.[1,2] • EM can transmit local data to LLM infrastructure, enhancing situational awareness.[3] ○ Pre-processing and feature extraction on edge devices can cut data transmission costs and time. ○ End devices may be able to provide fine-tuning and inference capabilities.

Potential Barrier: Are the necessary training and use-case data visible to the AI?	
Why It Matters	**What EM Should Know**
• Base LLMs are trained on generalized datasets. • Fine-tuning models post-training improves response accuracy and consistency plus model computational performance.[4] ○ Better for complex, multi-step, jargon-heavy, domain-specific, or high-consequence tasks. • Barriers like privacy, security, and proprietary information protections may block model access to fine-tuning data.	• EM tasks often require fine-tuned models trained on domain-specific data and/or RAG models that incorporate real-time data into outputs. • Situational awareness relies on multi-source data. • Incident data is often siloed due to disconnected agencies.[5] • Data silos block real-time decisions and can make recovery efforts costly and manual.

[1] Betzsold, N., Barr, J., Lesperance, A., Bartholomew, R., Ortega, S., Sleiman, C., Disney, M., & Tietje, G. (2024, June). Emergency Management of Tomorrow Research: emergency operations center of the future recommendations. *Pacific Northwest National Laboratory*. PNNL-38245.

[2] Betzsold, N., Otte, K., Barr, J., Lesperance, A., Bartholomew, R., Gray, J., Sleiman, C., Hagen, A., Disney, M., & Tietje, G. (2025, September). Emergency Management of Tomorrow Research: emergency operations center of the future technology exercises. *Pacific Northwest National Laboratory*. PNNL-36082

[3] Zhang, X., Xie, G., Huang, Y., Xiong, Z., Liu, J., Cui, S., Sun, S., & Sherman Shen, X. (2025). Edge intelligence in the Generative Artificial Intelligence Era. *IEEE Wireless Communications*, 1–9. https://doi.org/10.1109/mwc.2025.3599652

[4] Parthasarathy, V., Zafar, A., Khan, A., & Shahid, A. (2024). (tech.). *The Ultimate Guide to Fine-Tuning LLMs from Basics to Breakthroughs: An Exhaustive Review of Technologies, Research, Best Practices, Applied Research Challenges and Opportunities* (pp. 6–13). Dublin, Ireland.

[5] Lurie, R. & Stanko, K. (2024, October). The state of data in emergency management: challenges in integration and a path forward. Peregrine. https://media.erepublic.com/document/[2024]_IEM-FINAL-1001.pdf

Potential Barrier: Are the necessary training and use-case data visible to the AI?	
Why It Matters	**What EM Should Know**
• Horizontal data sharing happens across organizations, agencies, and public entities. • Vertical data sharing happens internally to organizations and across levels of government.[1] • Updated data sharing is necessary when information changes after the model's training and/or fine-tuning phases. ○ Retrieval augmented generation (RAG) models can use these up-to-date data to increase response relevancy.	• Historical disaster data is frequently siloed and underused for mitigation.[5] • Integrated tools and data-sharing can improve EM decision-making.[5]

4.2 Culture and Training

Culture indicates the willingness of users to adopt AI tools into their workflows, while training refers to users' knowledge, skills, and comfort in using AI effectively. Users must be able to understand and adhere to AI capabilities, limitations, and usage procedures.

Potential Barrier: Will users adapt to new AI technologies or remain in their comfort zone?	
Why It Matters	**What EM Should Know**
• Multimodal LLMs differ in usability, accessibility, and processing needs. • Some models are easier for non-technical users. • Immersive interfaces offer an interactive user experience but can cause motion sickness (among other usability impacts) and may face device limitations.[2] • Humans struggle to adapt when routines change. • Both humans and machines must improve adaptability in the AI era.[3]	• Staff demographics affect AI adoption.[4,5] • Adoption thrives on attention and free time.[6] • EM should foster exploration of AI tools. • AI must be integrated into daily tasks due to limited free time in EM. ○ Routine use builds AI tool familiarity.

[1] Betzsold, N., Otte, K., Barr, J., Lesperance, A., Bartholomew, R., Gray, J., Sleiman, C., Hagen, A., Disney, M., & Tietje, G. (2025, September). Emergency Management of Tomorrow Research: emergency operations center of the future technology exercises. *Pacific Northwest National Laboratory*. PNNL-36082.

[2] Bieniek, J., Rahouti, M., & Verma, D. (2024, November). Generative AI in multimodal user interfaces: trends, challenges, and cross-platform adaptability. https://doi.org/10.48550/arXiv.2411.10234

[3] Schmelzer, R. (2024, September 16). *The role adaptability plays in Generative AI*. Forbes. https://www.forbes.com/sites/ronschmelzer/2024/09/13/the-role-adaptability-plays-in-generative-ai/

[4] Morris, Michael G., and Viswanath Venkatesh. "Age differences in technology adoption decisions: Implications for a changing work force." *Personnel psychology* 53, no. 2 (2000): 375-403.

[5] Rojas-Mendez, Jose I., Ananthanarayanan Parasuraman, and Nicolas Papadopoulos. "Demographics, attitudes, and technology readiness: A cross-cultural analysis and model validation." *Marketing Intelligence & Planning* 35, no. 1 (2017): 18-39.

[6] Nagpal, G. K., & Mitra, A. (2025, November 5). *Research: When are customers willing to try a new technology?*. Harvard Business Review. https://hbr.org/2025/11/research-when-are-customers-willing-to-try-a-new-technology

Potential Barrier: Do users know when the AI tools should be used?	
Why It Matters	**What EM Should Know**

• Generative AI complements humans: AI is weaker in critical thinking and collaboration, while repetitive tasking can become challenging for humans.
• Effective human-machine teaming requires AI tools to achieve both high technological readiness levels (TRLs) and high human readiness levels (HRLs).[1]
• HRLs include ethical and practical guidelines, such as defining clear roles between the tool and the human.[1]
• Organizations should confirm:[1]
  ○ AI integrates clearly into workflows.
  ○ Humans know AI's role in various tasks.
  ○ Tasks between AI and humans are well-defined.
  ○ Current human workflows are understood.
  ○ Workflows are generalizable to more than one user.
	• The nature of EM work necessitates that AI tools have high HRLs prior to deployment.
	• EM users should clearly understand their tasks, objectives, and constraints without automation involved prior to incorporating AI.[1]
	• Workflow and requirements analyses can improve AI deployment outcomes and process improvement potential.[1]
	• Users should understand appropriate conditions of tool usage to ensure ethical and trustworthy deployment.[1]

Potential Barrier: Will users rely on the AI tool in high-stress situations?	
Why It Matters	**What EM Should Know**

• In emergencies, humans tend to rely on familiar tools.
• Stressful situations may lead users to avoid AI unless the users are already comfortable with it.
• Trust and understanding prevent reverting to old methods.[2]
• Infrequent events limit familiarity with AI tools used during said events.
	• EM practitioners want AI tools actively used rather than ignored.
	• Of EM organizations surveyed by Samsara:[3]
	  ○ 79% report that frontline teams lack training in digital tools and
	  ○ 82% cite growing data and technology training needs due to rising incident frequency.
	• Technology should be easy to use for all EM staff, not just experts.[4]
	• Ongoing training ensures effective AI use during incidents.

[1] Baweja, J. & Jefferson, B. (2024). *Human Factors in the Discovery Phase of TRLs and HRLs* [Conference presentation]. AIDA Palooza, Richland, WA, United States.

[2] Donahue, A., & Tuohy, R. (2006). Lessons we don't learn: a study of the lessons of disasters, why we repeat them, and how we can learn from them. *Homeland Security Affairs*, *2*.

[3] Government Fleet (2025, October). *New Samsara data reveals dangerous gap in disaster readiness*. Government Fleet – United States. https://www.government-fleet.com/10248923/new-samsara-data-reveals-dangerous-gap-in-disaster-readiness

[4] Pine, J. C. (2018). 1.2.2 Ease of Use of Technology. In *Technology and Emergency Management* (2nd ed., pp. 5–6). essay, Wiley.

Potential Barrier: Is an AI solution overly complex for the users' needs?	
Why It Matters	**What EM Should Know**
• Generative AI is not always the best solution to an organization's problems. • Overengineering unnecessarily creates overly complex systems. • Simpler options, like dashboards and organizational changes, may work better at times.[1] • Basic tools can be prerequisites for effective AI deployment.	• AI overhype can lead to poor deployment decisions and negative impacts. • Defining tasks and desired outcomes upfront can ensure AI is an appropriate solution. • Work domain analyses (among others) can identify gaps that are appropriate to address using AI.[2] • AI solutions should be evaluated ethically and empirically during development and requisitioning.[3]

Potential Barrier: Can users assess the quality of AI inputs and outputs?	
Why It Matters	**What EM Should Know**
• Assessing the quality of LLM inputs and outputs requires expertise in math, statistics, computer science, and data domains. • Prompt engineering – developing and rephrasing LLM instructions – improves response quality. • Skilled prompts elicit better reasoning and answers from models. • Training on concepts that indicate LLM response reliability can help users more effectively interact with models.	• EM staff need literacy in evolving technologies and their relevant and timely application to EM practice.[4] • Innovation awareness helps EM organizations assess security, potential utility, and required user expertise. • Recruiting and training should focus on cultivating technology skills. • Future EM teams may need to prioritize skills like prompt engineering and tech proficiency.

4.3 Resources

Resources pertain to the accessibility and reliability of the technological infrastructure, networking, power, and IT expertise that supports the deployment, integration, and effective use of AI.

[1] Myakala, P. K. (2024, November 22). *AI/ML hype-driven overengineering: The hidden cost of buzzword fever*. Medium. https://medium.com/@praveen.k.myakala/ai-ml-hype-driven-overengineering-the-hidden-cost-of-buzzword-fever-197fbd7e92b2

[2] Baweja, J. & Jefferson, B. (2024). *Human Factors in the Discovery Phase of TRLs and HRLs* [Conference presentation]. AIDA Palooza, Richland, WA, United States.

[3] Johnson, B., & Menzies, T. (2024). Ai over-hype: A dangerous threat (and how to fix it). *IEEE Software*, *41*(6), 131–138. https://doi.org/10.1109/ms.2024.3439138

[4] Betzsold, N., Barr, J., Lesperance, A., Bartholomew, R., Ortega, S., Sleiman, C., Disney, M., & Tietje, G. (2024, June). Emergency Management of Tomorrow Research: Emergency Operations Center of the Future Recommendations. *Pacific Northwest National Laboratory*. PNNL-38245.

Potential Barrier: Do existing systems support interoperability with new ones?	
Why It Matters	**What EM Should Know**
• Enterprise system complexity may block AI integration due to interface access (i.e., API availability), data formats, or cloud compatibility.[1] • Modernization boosts efficiency and enables AI integration. • Preparing source data and the user's queries before the main search (i.e., pre-retrieval) help legacy systems feed data into AI models for better contextual awareness and precision.[2]	• Current EOC information management systems can lack interoperability. • EM organizations should prioritize vendor-agnostic data systems and interoperability when choosing technologies.[4] • Existing software does not always automatically integrate with third-party generative AI applications, necessitating pre-retrieval mechanisms.

Potential Barrier: Is reliable network access available for data transfer?	
Why It Matters	**What EM Should Know**
• Generative AI requires massive datasets, which are costly and create data transmission delays.[3] • Transmission delays may reduce utility of LLM responses and other outputs. • Current cloud-based centralization can delay task completion and rely on outdated data, hindering user experience and increasing hallucinations.[4] • Some cloud-based AI models limit usage and number of responses per time period.[5] • Edge devices reduce transmission delays but may struggle with computational demands. • Generative AI strains EI significantly more than traditional AI models.[4]	• Close user proximity to generative AI models reduces latency and bandwidth consumption while increasing privacy protections, energy efficiency, and contextual awareness. • All phases of EM, but especially response, require (near) real-time information for critical decision-making. • EI can bridge the “last mile gap” between users and edge data.[3] • Model compression enables some generative AI tasking to take place on devices like smartphones. • Strong networking linking edge data, generative AI models, and users will be imperative for EM use cases.

[1] IT Convergence Blog. (2025, February 18). *Guide to integrating generative AI with your enterprise systems*. IT Convergence. https://www.itconvergence.com/blog/integrating-generative-ai-with-your-enterprise-systems/

[2] Du, C.-M., Tseng, C.-J., & Lu, C.-C. (2025). Bridging knowledge gaps: Leveraging legacy systems to enhance generative AI responses. *Intelligent Systems and Applications: Proceedings of the 2025 Intelligent Systems Conference (IntelliSys)*, *4*, 85–97. https://doi.org/10.1007/978-3-031-99965-9_6

[3] Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., & Zhang, J. (2019). Edge intelligence: Paving the last mile of artificial intelligence with Edge Computing. *Proceedings of the IEEE*, *107*(8), 1738–1762. https://doi.org/10.1109/jproc.2019.2918951

[4] Zhang, X., Xie, G., Huang, Y., Xiong, Z., Liu, J., Cui, S., Sun, S., & Sherman Shen, X. (2025). Edge intelligence in the Generative Artificial Intelligence Era. *IEEE Wireless Communications*, 1–9. https://doi.org/10.1109/mwc.2025.3599652

[5] Cao, Y., Li, S., Liu, Y., Yan, Z., Dai, Y., Yi, P., & Sun, L. (2023, March). A comprehensive survey of AI-generated content (AIGC): a history of generative AI from GAN to ChatGPT. https://doi.org/10.48550/arXiv.2303.04226

Potential Barrier: Do users have reliable access to power for devices using AI?	
Why It Matters	**What EM Should Know**
• Generative AI demands resource-intensive computations, lowering energy efficiency. • Edge devices like laptops and smartphones have limited resources. • Testing demonstrated that devices require up to 50% more battery runtime for effective AI usage.[1] • AI usage strains battery reliability, especially during power disruptions.[2]	• EOC operations and related incidents often involve power outages. • Backup power may fail – FEMA reports a 15% rise in generator failure rates after 24+ hours of continuous use.[3] • Backup power failures are further compounded by poor maintenance and testing. • Edge devices with power optimization will be vital as AI use grows in EM. • Adequate testing and maintenance of backup power systems should be implemented. • Using more computationally efficient AI models can alleviate some energy reliability concerns.

Potential Barrier: Do users have reliable access to IT and AI subject matter experts?	
Why It Matters	**What EM Should Know**
• Fine-tuning and other domain development efforts require generative AI modeling expertise.[4] • Deployment requires skills in software, cloud integration, and cybersecurity.[4] • Translating business and usage requirements needs judgement and contextualization.[4] • Model maintenance and troubleshooting may need to happen within end-user organizations. • Legal and governance roles will emerge to ensure ethical AI use.[4]	• EM practitioners must learn new skills. • EM organizations must hire staff with AI expertise. • Continual training is essential to ensure staff can troubleshoot AI and keep pace with rapid industry developments. • Early engagement can help build a technology-literate workforce.[5]

[1] Talluri, R. (2025, July 24). *The AI power drain: Why battery limitations threaten the future of Mobile AI*. Enovix. https://www.enovix.com/the-ai-power-drain-why-battery-limitations-threaten-the-future-of-mobile-ai/

[2] Daud, A., Al Abdouli, K. M., & Badshah, A. (2025). Emerging computing tools for emergency management: Applications, limitations and future prospects. *IEEE Open Journal of the Computer Society*, *6*, 627–644. https://doi.org/10.1109/ojcs.2025.3563759

[3] Federal Emergency Management Agency. (2025). (rep.). *Power Outage Incident Annex to the Response and Recovery Federal Interagency Operational Plan*. Washington D.C.: U.S. Department of Homeland Security.

[4] Relyea, C., Maor, D., Durth, S., & Bouly, J. (2024, August 7). *Gen Ai's next inflection point: From employee experimentation to organizational transformation*. McKinsey & Company. https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/gen-ais-next-inflection-point-from-employee-experimentation-to-organizational-transformation

[5] Betzsold, N., Barr, J., Lesperance, A., Bartholomew, R., Ortega, S., Sleiman, C., Disney, M., & Tietje, G. (2024, June). Emergency Management of Tomorrow Research: emergency operations center of the future recommendations. *Pacific Northwest National Laboratory*. PNNL-38245.

Potential Barrier: Do users have reliable, up-to-date hardware and other computing resources?	
Why It Matters	**What EM Should Know**
• Local AI processing will soon be key for expected model responsiveness.[1] • AI is being increasingly embedded in edge devices. ○ 30% of smartphones and 43% of PCs delivered in 2025 are predicted to be AI-enabled.[1] • Generative AI computational demands exceed traditional edge device hardware capabilities. • User interfaces must evolve with changing user expectations. • User interface questions include ideal design, app-specific interfaces, and immersive interface impacts.[2]	• EM edge devices like laptops, smartphones, cameras, sensors, and drones will likely need upgrades for emerging AI capabilities. • Direct model integration into devices will grow, increasing local processing demands. • Non-traditional user interfaces like virtual reality and voice control may improve EM user interactions with AI.[2] • Without hardware advancements, EM use of generative AI will be limited.

4.4 Trust and Explainability

Trust indicates the user's confidence in the accuracy and reliability of AI inputs and outputs. Explainability pertains to the ability of AI models to disclose their reasoning to users, enabling informed trust evaluations. Users must be able to understand when they can and cannot trust an AI model in order to incorporate the model into decision-making.

Potential Barrier: Do users know when to trust AI inputs?	
Why It Matters	**What EM Should Know**
• Techniques like fine-tuning and RAG can enhance generative AI performance and output. • These techniques require high-quality, domain-specific data. • Assessing data quality for LLMs is multi-faceted and complex. • Poor data causes model response inconsistencies, noncompliance, and hallucinations.[3] • Trustworthy use case data is vital for certain LLM tasks.	• EM practitioners often worry about data quality, especially during the response phase. • EM work is often sensitive and high consequence, which means it is imperative that AI responses be trustworthy. • Poor data hampers decision-making using AI responses.[4] • Damaged sensors can cause missing or inconsistent data.

[1] Talluri, R. (2025, July 24). *The AI power drain: Why battery limitations threaten the future of Mobile AI*. Enovix. https://www.enovix.com/the-ai-power-drain-why-battery-limitations-threaten-the-future-of-mobile-ai/

[2] Bieniek, J., Rahouti, M., & Verma, D. (2024, November). Generative AI in multimodal user interfaces: trends, challenges, and cross-platform adaptability. https://doi.org/10.48550/arXiv.2411.10234

[3] Yu, X., Zhang, Z., Niu, F., Hu, X., Xia, X., & Grundy, J. (2024). What makes a high-quality training dataset for large language models: A practitioners' perspective. *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*, 656–668. https://doi.org/10.1145/3691620.3695061

[4] Jayawardene, V., Huggins, T. J., Prasanna, R., & Fakhruddin, B. (2021). The role of data and information quality during disaster response decision-making. *Progress in Disaster Science*, *12*. https://doi.org/10.1016/j.pdisas.2021.100202

Potential Barrier: Do users know when to trust AI inputs?	
Why It Matters	**What EM Should Know**
• Understanding input data quality is key to understanding the reliability of generative AI results.	• Improvement recommendations are extensive but often challenging to implement.[1] • LLMs may need models to assess input data quality in their pipelines. • Automated alerts or stopping points can notify users when there is poor output quality or when the system is unable to produce useful results due to unreliable or incorrect input data.

Potential Barrier: Do users know when to trust AI outputs?	
Why It Matters	**What EM Should Know**

• AI misuse examples include fake books in reading lists and fabricated court cases in submitted precedent briefs, among others.[2]
• Trust calibration is critical for confident and reasonable decision-making based on generative AI outputs.
• Users must continuously adapt their understanding of a model's reliability during usage.
• Both under- and overconfidence in AI outputs can lead to serious risks.[3]
  ○ Under-confidence may cause reduced tool use or human-machine teaming breakdowns.
  ○ Overconfidence may cause harmful or suboptimal decisions.
• Early errors in LLMs are shown to significantly reduce human trust and subvert trust rebuilding.[4]
	• Mistakes by EM organizations can erode long-term public trust, and EM concerns about public mistrust make EM practitioners hesitant to adopt AI.[5]
	• LLMs will inherently make mistakes, requiring strategies to help users calibrate trust.
	• Confidence scores and reasoning explanation can improve model explainability and trust.[4]
	• Guidelines can help users understand when to trust generative AI outputs.
	• EMOTR tabletop exercise participants suggest treating AI as an intern rather than an assistant and review all work prior to decision-making.
	• EM practitioners prefer AI to support, not replace, human decision-making.[5]

[1] Betzsold, N., Barr, J., Lesperance, A., Bartholomew, R., Ortega, S., Sleiman, C., Disney, M., & Tietje, G. (2024, June). Emergency Management of Tomorrow Research: emergency operations center of the future recommendations. *Pacific Northwest National Laboratory*. PNNL-38245.

[2] Olavsrud, T. (2025, August 7). *11 famous AI Disasters*. CIO. https://www.cio.com/article/190888/5-famous-analytics-and-ai-disasters.html

[3] Lebiere, C., Blaha, L. M., Fallon, C. K., & Jefferson, B. (2021). Adaptive cognitive mechanisms to maintain calibrated trust and Reliance in automation. *Frontiers in Robotics and AI*, *8*. https://doi.org/10.3389/frobt.2021.652776

[4] Martell, M. J., Baweja, J. A., & Dreslin, B. D. (2024). Mitigative strategies for recovering from Large Language Model Trust violations. *Journal of Cognitive Engineering and Decision Making*, *19*(1), 76–95. https://doi.org/10.1177/15553434241303577

[5] Betzsold, N., Otte, K., Barr, J., Lesperance, A., Bartholomew, R., Gray, J., Sleiman, C., Hagen, A., Disney, M., & Tietje, G. (2025, September). Emergency Management of Tomorrow Research: emergency operations center of the future technology exercises. *Pacific Northwest National Laboratory*. PNNL-36082.

Potential Barrier: Are guardrails in place to prevent misuse of AI tools?	
Why It Matters	**What EM Should Know**
• Generative AI aids task automation, data analysis, and intelligence generation. • Negative impacts on humans include inappropriate decision biasing, reduced awareness, reduced situational understanding, and poor performance. • AI models may unintentionally propagate biases or enable fraud. • Guardrails are needed to protect users from harm.[1]	• EM practitioners should ensure all traditional and generative AI models follow recognized safety guidelines for critical tasks.[1] • Human factors research can help improve AI safety and decision-making.[1] • Fail-safes, alerts, and breakpoints can help prevent misuse by humans or automation. • Users should be trained on AI limits, failures, and alternative methods.[1]

[1] Human Factors and Ergonomics Society (N.d.). (publication). *Artificial Intelligence Guardrails for Human Use*. Retrieved December 3, 2025, from https://www.hfes.org/Portals/0/HFES%20Policy%20Statement%20-%20AI%20guardrails%20for%20human%20use_-Rev%207.pdf?ver=g5kPVsH-yzg70SrlW1Pl5w%3d%3d.

5.0 Promising Technologies and Approaches

This section details the specific technologies and approaches through which EM-relevant data and information can be identified, acquired, owned, pre-processed, and made usable by AI systems for EM use cases (Figure 4). Additionally, each section details the pertinent EM functionality of the layer, followed by the set of specific technology profiles that describe both the technologies and stakeholder-specific relevance and potential opportunities. Detailed profiles are available in Appendix A.

Layer of AI Stack	Featured Technologies and Approaches
1. Infrastructure	• Data storage architectures • Data access architectures
2. Agent Internet	• Autonomous systems • Agent schema standards
3. Protocols	• Agent-to-agent protocol • Model context protocol • Agent gateways • Data interoperability standards • Data provenance and trust frameworks
4. Tooling and Enrichment	• Data discovery for improved accuracy: • Fine tuning • Retrieval-Augmented Generation • Data quality and safety controls for AI systems • Privacy-preserving collaboration and data sharing • Real-time intelligence and continuous model enhancement
6. Memory and Personalization	• Handling missing and/or degraded data quality • Handling missing data modalities • Controlling protocols for sudden retroactive access to data • Multi-jurisdictional knowledge access and compliance • Contextual knowledge integration for enhanced decision support
8. Operations and Governance	• Real-time data ingestion services • Data sovereignty and cross-jurisdictional compliance • Accountability, audit, and incident response for AI systems • Trust and quality assurance for AI operations

Figure 4. AI stack layers from the framework analyzed to support AI advancement and adoption in EM. Layer 5 (cognition and reasoning) and layer 7 (applications) are absent as they are out of scope for the objectives of this effort.

5.1 Infrastructure (Layer 1)

The Infrastructure Layer of the AI stack encompasses foundational technologies that manage how organizations store and access data to support AI capabilities. Data storage architectures (e.g., data warehouses, data lakes, and data lakehouses [see Appendix A, Figure A.1 for details]) serve distinct roles and complement each other in enabling scalable, reliable, and flexible data management. A data warehouse provides structured, cleaned data for repeatable analysis and reporting, while a data lake captures raw information for exploratory analytics, including AI model training. The data lakehouse blends the strengths of both, offering scalable analytics with accuracy and flexibility. Together, these approaches help organizations achieve situational awareness, compliance, and governance while adapting to varying data demands. Additionally, data access architectures like data fabric and data mesh enhance the discoverability, accountability, and governance of data across organizations. A data fabric connects data systems and ensures secure access through catalogs, lineage tracking, and policy enforcement. A data mesh assigns ownership of well-defined data products to domain experts, ensuring accuracy, clarity, and faster delivery. When combined, these architectures improve collaboration, AI reliability, and the safe integration of data across agencies.

For EM, employing these architectures supports robust situational awareness, auditable reports, real-time forecasting, and scalable cross-agency analytics. However, adoption challenges include integrating scattered legacy systems, inconsistent data definitions, siloed information, and limited metadata practices.

To address these gaps, state and local agencies should inventory data, implement shared catalogs, enforce data governance policies, and pilot AI tools to ensure reproducibility and reliability. Policymakers and vendors must support interoperability, security, open standards, and modernization funding to advance AI readiness. Ultimately, achieving reliable data infrastructure and access across EM builds a strong foundation for actionable, explainable AI.

Infrastructure (Layer 1)

Infrastructure technology profiles can be found in Appendix A (Section A.1) and include the following profiles:

- Data Storage Architectures
- AI Data Access Architectures

5.2 Agent Internet (Layer 2)

The Agent Internet Layer establishes the universal connectivity fabric through which agents discover, access, and interact with external data sources, services, and one another. It abstracts transport, authentication, and service-discovery concerns so agents can operate across heterogeneous environments. Through natural language inputs and integration with systems like Computer-Aided Dispatch (CAD), autonomous systems can be activated with minimal human intervention. Their functionality relies on advanced modeling, decision-making processes, and the ability to interpret unforeseen circumstances. Common use cases include aerial drones for search-and-rescue, hazardous material (HAZMAT) detection, and disaster damage assessments, emphasizing their ability to operate in dangerous or inaccessible environments.

Despite their promise, autonomous systems face critical challenges when applied to life-safety situations, including environmental awareness, reliability, and processing unexpected scenarios. These limitations underscore the need for rigorous testing, effective training data, and clear operational policies to ensure their successful deployment. As their capabilities evolve, emergency managers must benchmark their performance within narrowly defined scenarios before broader adoption.

To address these challenges, EM organizations can focus on identifying tasks where autonomous systems provide measurable benefits, establish key performance indicators, and collaborate with developers to refine technology. Standards and policymakers can take action to prioritize public safety datasets, standardized testing, liability frameworks, and funding mechanisms to enable reliable use in emergency operations. Vendors must clearly document their systems' capabilities and limitations while partnering with emergency organizations to bridge gaps in operational requirements and AI understanding. For AI leaders, ensuring mission-critical reliability, developing resilience for split-second decision-making in unforeseen scenarios, and implementing robust fail-safe mechanisms will be key to advancing autonomous systems tailored to public safety needs.

Agent Internet (Layer 2)

Agent internet technology profiles can be found in Appendix A (Section A.2) and include the following profiles:

- Autonomous Systems
- Agent Schema Standards

5.3 Protocols (Layer 3)

The Protocol Layer of the AI stack focuses on enhancing EM systems through interoperability, security, provenance, and standardized communication frameworks. A2A protocols streamline coordination across organizations for tasks like wildfire response and resource allocation, while hierarchical systems align with the Incident Command System (ICS) to distribute strategic and tactical activities efficiently. Agent gateways ensure secure interactions with authentication, Role-Based Access Control, and dynamic routing for resiliency. MCP simplifies AI integration with platforms like databases, sensors, and Geographic Information Systems (GIS), enabling seamless real-time data flow for resource allocation, predictive analytics, and threat detection. Data Interoperability Standards (such as National Information Exchange Model (NIEM) and open formats) prevent silos and promote multi-agency coordination, empowering AI systems to analyze and interpret data consistently. Provenance and trust frameworks like Coalition for Content Provenance and Authenticity (C2PA) and Software Bill of Materials (SBOM) provide traceability for AI-generated outputs, verifying content authenticity, documenting training data, and ensuring accountability in life-critical decisions, while enhancing public trust and supporting after-action reviews and legal defenses.

To address these challenges, emergency managers, policymakers, vendors, and AI leaders must collaborate to effectively implement these protocols, enabling secure, interoperable, and trusted AI systems. Actions include creating dashboards for monitoring agent-to-agent communications, supporting interoperability standards in RFPs, establishing liability frameworks for AI decisions, and embedding provenance tracking in all AI systems to ensure transparency and accountability. By prioritizing these standards and technologies, the protocol layer solidifies the foundation for scalable and reliable AI integration in EM operations.

Protocols (Layer 3)

Protocols technology profiles can be found in Appendix A (Section A.3) and include the following profiles:

- Agent to Agent
- Agent Gateway
- Model Context Protocol
- Data Interoperability Standards
- Data Provenance and Trust Frameworks

5.4 Tooling and Enrichment (Layer 4)

The Tooling and Enrichment Layer of the AI stack enhances EM by optimizing data processing, model accuracy, and operational reliability. Technologies like Fine-Tuning customize general language models for EM-specific tasks such as risk analysis and evacuation planning, leveraging domain-specific datasets to improve relevance, accuracy, and privacy while mitigating risks like hallucinations. RAG integrates up-to-date data for real-time decisions by efficiently organizing information in vector databases, with Retrieval-Augmented Fine-Tuning combining fine-tuning and RAG to balance domain knowledge and adaptability. Data Quality and Safety Controls, including content moderation, anomaly detection, and personally identifiable information (PII) redaction, safeguard models from unreliable data, and privacy risks during incidents. Privacy-Preserving Collaboration, through tools like data clean rooms and federated learning, enables secure multi-agency coordination by allowing joint analysis without exposing sensitive data. Real-Time Intelligence and Continuous Model Enhancement further empower AI systems to adapt to evolving emergencies, with streaming pipelines and time-series feature engineering providing low-latency, predictive capabilities to inform dynamic decision-making and issue early warnings for hazards.

Adoption recommendations include defining clear scopes for AI deployment, fine-tuning models with high-quality datasets, establishing data sharing agreements, and implementing robust safeguards for data quality and privacy. Policymakers should mandate interoperability, privacy-preserving standards, and liability frameworks. Vendors should focus on embedding explainability, real-time intelligence, and iterative model improvement into their solutions while enabling cross-agency collaboration. AI leaders must ensure systems are resilient, auditable, and capable of real-time adaptation to EM-specific demands, ultimately facilitating safer, faster, and more reliable emergency responses.

Tooling and Enrichment (Layer 4)

Tooling and enrichment technology profiles can be found in Appendix A (Section A.4) and include the following profiles:

- Data Discovery for Improved Accuracy: Fine-Tuning
- Data Discovery for Improved Accuracy: Retrieval-Augmented Generation
- Data Quality and Safety Controls for AI Systems
- Privacy-Preserving Collaboration and Data Sharing
- Real-Time Intelligence and Continuous Model Enhancement

5.5 Memory and Personalization (Layer 6)

The Memory and Personalization Layer of the AI stack focuses on improving the resilience and adaptability of AI systems in EM scenarios by addressing challenges with data quality, missing data modalities, retroactive data access, multi-jurisdictional compliance, and contextual knowledge integration. EM frequently faces degraded or incomplete data streams during disasters due to damaged sensors or disrupted infrastructure, impacting AI decision support accuracy.

To address these challenges, EM stakeholders can employ data enhancement engines that mitigate these concerns by cleaning, imputing, and analyzing data dependencies, ensuring models remain functional with uncertain inputs. When specific data modalities become unavailable, AI systems can rely on redundant or proxy streams, enabled by correlation analysis, causal inference, or creative repurposing of resources like drones. Retroactive access to large amounts of data following connectivity restoration poses another challenge, but systems can backfill and recalibrate datasets, ensuring critical information gaps are addressed. Federated RAG and jurisdiction-partitioned indices empower AI systems to retrieve knowledge while respecting privacy, sovereignty, and regulatory boundaries, allowing multi-agency coordination without centralizing sensitive data. Lastly, hybrid knowledge architectures combining vector databases and knowledge graphs enhance recommendations by integrating semantic relevance with operational constraints, ensuring decisions are both contextually appropriate and actionable. These capabilities transform AI into robust, adaptive decision support systems that can navigate the complex, dynamic, and multi-jurisdictional demands of EM operations.

Memory and Personalization (Layer 6)

Memory and personalization technology profiles can be found in Appendix A (Section A.5) and include the following profiles:

- Handling Missing and/or Degraded Data Quality
- Handling Missing Data Modalities
- Handling Sudden Retroactive Access to Data
- Multi-Jurisdictional Knowledge Access and Compliance
- Contextual Knowledge Integration for Enhanced Decision Support

5.6 Governance and Operations (Layer 8)

The Governance and Operations Layer is focused on providing integration, trust, and compliance for AI systems. Key measures include real-time data ingestion services to rapidly incorporate new data streams like on-scene responder sensors, drones, and online discourse, facilitated by formatting standards and LLM-generated adapters. Data Sovereignty and Compliance uses residency routing, egress controls, and data loss prevention systems to protect sensitive information while meeting jurisdictional and regulatory requirements. Control/data plane separation enables centralized coordination without sacrificing data sovereignty. Accountability and Incident Response ensure traceability through audit logging, retention policies, and kill switches for quick system disabling during failures or cyberattacks. Trust and Quality Assurance maintains reliability via continuous data quality monitoring, drift detection, model registries, and continuous integration/continuous deployment (CI/CD) pipelines for safe deployment of updates. These measures bolster operational reliability, compliance, and adaptability for EM organizations during dynamic disaster scenarios.

To address governance and operations challenges, EM organizations can implement Real-Time Data Ingestion Services to adapt quickly to new data streams using LLM-based adapters and standardized formats. Policymakers can focus efforts on developing consistent non-technical data-sharing protocols to enhance cross-jurisdictional collaboration. Robust Data Sovereignty and Compliance systems should enforce privacy through routing, controls, and compliance auditing while allowing regional coordination. EM vendors can embed accountability with audit logging, e-discovery tools, retention policies, and kill switches to safeguard operations during crises. Finally, trust and quality assurance protocols should be incorporated into these systems to leverage continuous monitoring, drift detection, model registries, and systematic testing.

Governance and Operations (Layer 8)

Governance and operations technology profiles can be found in Appendix A (Section A.6) and include the following profiles:

- Real-time Data Ingestion Services
- Data Sovereignty and Cross-Jurisdictional Compliance
- Accountability, Audit, and Incident Response for AI Systems
- Trust and Quality Assurance for AI Operations

6.0 Summary: Key Opportunities for EM Stakeholders

6.1 EM Community

EM organizations operate in high-stakes environments where timely, informed, and collaborative decision-making is critical, and AI systems offer transformative opportunities to enhance situational awareness, resource allocation, and operational coordination. To fully leverage these capabilities, organizations must prioritize seamless data access and integration across data streams, including real-time feeds from sensors, GIS, dispatch systems, drones, and online discourse, ensuring interoperability and reliable data sharing across agencies. Data quality and reliability are equally critical, as disasters can degrade or fragment data, requiring robust validation, enhancement, and traceability frameworks to maintain accurate and trustworthy AI outputs. Privacy and security challenges must also be addressed with technologies like privacy-preserving analytics and on-device preprocessing safeguarding sensitive information while enabling collaboration across jurisdictions. Additionally, AI tools such as autonomous systems, RAG, and fine-tuned models provide powerful capabilities but require human oversight and awareness of data dependency limitations to ensure effective deployment. Finally, real-time intelligence and adaptive systems are crucial for responding to dynamic disaster scenarios, with AI continuously updating recommendations to reflect evolving conditions and challenges. By addressing these interconnected priorities, EM organizations can maximize the potential of AI systems, improving disaster response efficiency, protecting resources, and ultimately saving lives.

EM Community Opportunities

- Enable Data Management and Governance
- Implement Technology and Training
- Mitigate Risks and Enhance Resilience
- Scale and Pilot AI Solutions
- Promote Collaboration

The following activities can prepare the EM community to implement EM-enhancing AI technologies (see technology profiles in Appendix A for additional details):

1. **Enable Data Management and Governance:**
    a. Inventory and classify high-value datasets by sensitivity, update frequency, and importance.
    b. Establish secure data exchange agreements, interoperability standards, and clear policies for data sharing across agencies.
2. **Implement Technology and Training:**
    a. Deploy systems like privacy-preserving analytics, dashboards, provenance frameworks, and AI-enhanced tools for improved situational awareness and faster decision-making.
    b. Train personnel in stewardship, querying tools, data workflows, and the interpretation of AI recommendations.
3. **Mitigate Risks and Enhance Resilience:**

a. Prepare for degraded or missing data by implementing redundant data streams, proxy modalities, and fallback mechanisms such as human validation or manual workflows.

b. Test incident response capabilities regularly (e.g., kill switch protocols, manual fallbacks, anomaly detection) to quickly address system failures or adversarial attacks.

4. **Scale and Pilot AI Solutions:**

   a. Pilot AI tools in low-risk scenarios to assess reliability before scaling up to multi-agency coordination or high-stakes tasks.

   b. Continuously update systems (e.g., hybrid knowledge architectures, CI/CD pipelines) to ensure flexibility and relevance in real-world deployment.

5. **Promote Collaboration:**

   a. Leverage federated systems for multi-jurisdictional coordination, ensuring shared analysis and mutual aid without centralizing sensitive data.

   b. Establish partnerships with public and private sectors to expand resources, improve data sharing, and enable better disaster recovery plans.

6.2 EM Policymakers/Standards

Policymakers and standards developers play a vital role in enabling effective AI adoption in EM by addressing key priorities such as interoperability, governance, accountability, and ethics. Seamless collaboration across agencies and jurisdictions requires policies that prioritize interoperability and data sharing through open formats that enhance accessibility and compatibility, breaking down barriers to multi-agency and cross-sector cooperation. Safeguarding sensitive information like PII and organizational intelligence is critical, demanding compliance with privacy laws, data sovereignty requirements, and governance standards that include transparent audit trails and provenance systems to ensure accountability in life-critical scenarios. Trust and accountability must be reinforced through mechanisms for audit logging, compliance verification, and clear liability frameworks, ensuring AI-driven outcomes are traceable and reliable. Resilient AI systems must be guided by standards that address real-time reliability and data quality, equipping systems to handle uncertainty, degraded inputs, and missing modalities with effective fallback mechanisms. Coordination across jurisdictions requires federated architectures that respect legal boundaries while promoting secure and efficient multi-agency collaboration. Additionally, ethical considerations must take center stage, with standards designed to minimize bias propagation in training data and AI outputs to ensure fairness and equity. Together, these priorities build a solid foundation for AI in EM, supporting systems that are interoperable, secure, reliable, and ethical.

Policymaker and Standards Opportunities

- Develop Standards for Interoperability and Collaboration
- Enhance Privacy and Security Safeguards
- Mandate Accountability and Governance Policies
- Support Resilience and Reliability
- Promote Coordination Across Jurisdictions
- Focus on Training, Awareness, and Bias Mitigation

The following activities can prepare the EM policymakers and standards developers to implement EM-enhancing AI technologies (see technology profiles in Appendix A for additional details):

1. **Develop Standards for Interoperability and Collaboration:**
    a. Implement open data formats and shared metadata protocols for effective cross-agency and multi-jurisdictional information sharing.
    b. Create certification programs ensuring vendors comply with standards for interoperability, provenance, and data quality, enabling trustworthy AI integration.
2. **Enhance Privacy and Security Safeguards:**
    a. Establish frameworks addressing PII redaction, data governance, and secure preprocessing at the source to protect sensitive information while promoting collaboration.
    b. Define audit trail requirements beyond traditional rules, ensuring AI systems produce reliable accountability measures compatible with privacy mandates.
3. **Mandate Accountability and Governance Policies:**
    a. Create policies requiring organizations to log AI operations comprehensively (e.g., data accessed, recommendations made, confidence scores).
    b. Clarify liability protections for organizations that follow documented governance and data handling standards while setting accountability benchmarks for poorly implemented systems.
4. **Support Resilience and Reliability:**
    a. Define real-time AI performance standards, including acceptable thresholds for data quality and system latency. Require systems to maintain operational capability during degraded conditions.
    b. Fund pilot programs to refine data quality protocols and test reliability standards in simulated environments.
5. **Promote Coordination Across Jurisdictions:**
    a. Develop frameworks for federated architectures that enable secure knowledge transfer across agencies while maintaining regional sovereignty over data.
    b. Create templates for multi-jurisdictional data-sharing agreements that align with legal and operational requirements, reducing friction during emergencies.
6. **Focus on Training, Awareness, and Bias Mitigation:**
    a. Establish ethical guidelines for dataset creation, sharing, and use to minimize bias propagation in AI outcomes.
    b. Train staff on compliance, accountability mandates, dataset stewardship, and the implications of AI decision-making.

6.3 EM Vendors

EM vendors play a crucial role in addressing fundamental challenges encountered in EM, including fragmented systems, data quality, scalability, and multi-agency collaboration to enhance disaster response and coordination. Effective and efficient data integration and

interoperability are vital, requiring solutions compatible with existing systems like CAD, Resource Management Systems (RMS), GIS, sensors, and messaging platforms, as well as tools that support open formats and standards such as NIEM for reliable cross-agency data sharing. To ensure data quality, security, and provenance, vendors must implement robust validation processes, content moderation, and PII redaction, while providing traceability and accountability through tamper-evident audit logs and standards like C2PA.

Building adaptable and scalable systems is essential, with real-time data ingestion, modular capabilities, and cross-jurisdictional collaboration that respond effectively to dynamic disaster environments. Vendors must also meet accountability and governance compliance requirements by embedding features like audit logging, data retention policies, compliance dashboards, and rapid response mechanisms such as kill switches to maintain operational integrity in emergencies. Furthermore, privacy-preserving and cross-agency collaboration technologies, including federated architectures, on-device preprocessing, and data clean rooms, are critical for safeguarding sensitive information while enabling seamless multi-jurisdictional cooperation. By addressing these priorities, EM vendors can deliver secure, interoperable, adaptable, and compliance-ready solutions that meet the critical needs of EM operations.

EM Vendor Opportunities

- Design Solutions Enabling Integration and Interoperability
- Implement Data Quality and Provenance Features
- Build Adaptive, Scalable Models
- Embed Accountability and Governance Tools
- Facilitate Privacy-Preserving Collaboration
- Build Trust via Explainability and Partnerships

The following activities can prepare the EM policymakers and standards developers to implement EM-enhancing AI technologies (see technology profiles in Appendix A for additional details):

1. **Design Solutions Enabling Integration and Interoperability:**
    a. Build systems that integrate streaming and batch data from multiple sources, exposing rich metadata, lineage, and granular access controls.
    b. Prioritize support for open standards like NIEM and develop APIs and software development kits (SDKs) that facilitate integration with legacy and emerging tools in multi-agency environments.
2. **Implement Data Quality and Provenance Features:**
    a. Develop robust data validation pipelines to flag anomalies, missing information, and biased data before it influences recommendations.
    b. Offer provenance capabilities, including audit trails tracking AI-generated inputs, outputs, and decision processes to support accountability requirements and after-action reviews.
3. **Build Adaptive, Scalable Models:**
    a. Create flexible systems that can rapidly incorporate new data streams, such as real-time sensor inputs, through mechanisms like LLM-based data adapters.

   b. Deploy solutions that scale to meet variable conditions with advanced architecture features like anomaly detection, time-series engineering, and hierarchical knowledge retrieval.

4. **Embed Accountability and Governance Tools:**
   a. Implement granular kill switches allowing selective shutdown of flawed AI components during emergencies to prevent cascading failures.
   b. Create compliance dashboards and e-discovery tools to help EM organizations manage documentation, demonstrate audit readiness, and satisfy regulatory mandates.
5. **Facilitate Privacy-Preserving Collaboration:**
   a. Incorporate federated AI setups and privacy-preserving measures like encrypted local data processing into system designs to enable multi-agency workflows while safeguarding sensitive information.
   b. Support secure cross-agency collaboration with transparent documentation, compliance audits, and regional infrastructure for federated retrieval and knowledge sharing.
6. **Build Trust via Explainability and Partnerships:**
   a. Develop explainability features in AI tools to help EM clients interpret model recommendations, trace data sources, and assess the reliability of outputs during high-pressure scenarios.
   b. Partner with standards organizations, AI developers, and complementary vendors to create pre-integrated solutions that meet EM-specific requirements seamlessly.

6.4 AI Leaders

AI leaders are foundational in developing AI-enabling solutions that address the complex challenges of EM, where high-stakes operations demand reliable, scalable, and transparent systems. AI tools must process high-velocity and data from sources like sensors, GIS platforms, and communication networks while adapting to legacy infrastructure, multi-agency coordination, and variable network conditions. These systems must align with public safety rules, terminology, and decision-making processes, gracefully handling unexpected scenarios to prevent increased risks during emergencies. Ensuring data quality and reliability is critical, as real-time data often includes missing entries, biases, and degraded quality.

AI systems must employ rigorous validation, anomaly detection, drift monitoring, and provenance tracking to maintain transparency and accountability. Interoperability and standards are essential for seamless integration across platforms and jurisdictions, with frameworks like NIEM and hybrid architectures combining vector databases and knowledge graphs to improve operational feasibility and cross-agency collaboration. At the same time, privacy-preserving technologies like federated learning and secure multi-party computation protect sensitive data and enable secure collaboration across organizations while mitigating privacy risks and adversarial threats. These systems must be highly resilient and scalable to operate reliably under challenging conditions such as missing data, degraded infrastructure, or sudden influxes of new information, incorporating fallback mechanisms, graceful degradation, and human-in-the-loop processes for critical decisions. By addressing these priorities, AI leaders can develop robust, adaptable, and secure solutions that advance EM operations and improve safety outcomes.

AI Leader Opportunities

- Develop and Test AI Systems for Emergency-Specific Scenarios
- Implement Robust Data Quality and Provenance Mechanisms
- Focus on Interoperability and Standards Support
- Embed Privacy-Preserving Collaboration
- Design Resilient, Scalable AI Systems
- Promote Transparency and Accountability

The following activities can prepare the AI Leaders to implement EM-enhancing AI technologies (see technology profiles in Appendix A for additional details):

1. **Develop and Test AI Systems for Emergency-Specific Scenarios:**
   a. Pilot low-risk workflows with EM organizations to assess system reliability and align solutions with operational needs such as real-time alerts, risk detection, and resource optimization.
   b. Build disaster-focused autonomous systems and AI agent-to-agent systems that meet mission-critical reliability and security standards, designed to complement, not replace, human decision-making.
2. **Implement Robust Data Quality and Provenance Mechanisms:**
   a. Establish data pipelines for continuous validation, anomaly detection, and enrichment to ensure data reliability under challenging conditions.
   b. Integrate provenance tracking and explainability interfaces into AI systems so EM practitioners can trace recommendations, verify authenticity, and understand both the confidence and limitations of outputs.
3. **Focus on Interoperability and Standards Support:**
   a. Design AI systems with native compatibility for standards like NIEM and adopt protocols like Parquet, Delta, and Iceberg for efficient batch processing and data sharing.
   b. Participate in standards working groups with EM organizations to align AI capabilities with domain-specific requirements and ensure interoperability across platforms.
4. **Embed Privacy-Preserving Collaboration:**
   a. Develop AI systems leveraging privacy-preserving approaches such as federated learning and clean rooms to enable secure collaboration across agencies while protecting sensitive data.
   b. Create tools that account for privacy and security requirements when sharing data with external partners, enforcing compliance while fostering operational alignment.
5. **Design Resilient, Scalable AI Systems:**
   a. Build AI systems capable of handling uncertainty, sudden data shifts, and dynamic operations. Provide mechanisms like fallback modes, graceful failure, and real-time adaptability to maintain reliable performance under stress.
   b. Establish CI/CD pipelines to test new features in simulated environments replicating emergency conditions, ensuring updates deploy safely and reliably during crises.
6. **Promote Transparency and Accountability:**

a. Create detailed audit logs and retention systems that document data provenance, model operations, and decision-making processes for compliance, accountability, and after-action reviews.
b. Ensure systems can clearly explain outputs alongside traceable metadata such as confidence scores, data lineage, and decision alternatives for transparency and trust.

Appendix A – Technology Profiles

A.1 Infrastructure

Infrastructure provides the foundational compute, storage, and network resources (on-premises, cloud, or edge) that power every other layer. Its purpose is to deliver resilient, scalable, and cost-optimized hardware and core platform services without prescribing a particular vendor or technology.

A.1.1 Data Storage Architectures

Data storage architectures describe how an organization keeps and organizes its information. Good storage architecture makes information easier to find, more reliable to use, and ready for both traditional reporting and modern AI. It balances structure and flexibility to provide needed insights today and adapt to new requests in the future. Three widely used storage architectures are the data warehouse, data lake, and data lakehouse (see Figure A.1). Each serves a different purpose, and most organizations use them together in a complementary fashion.

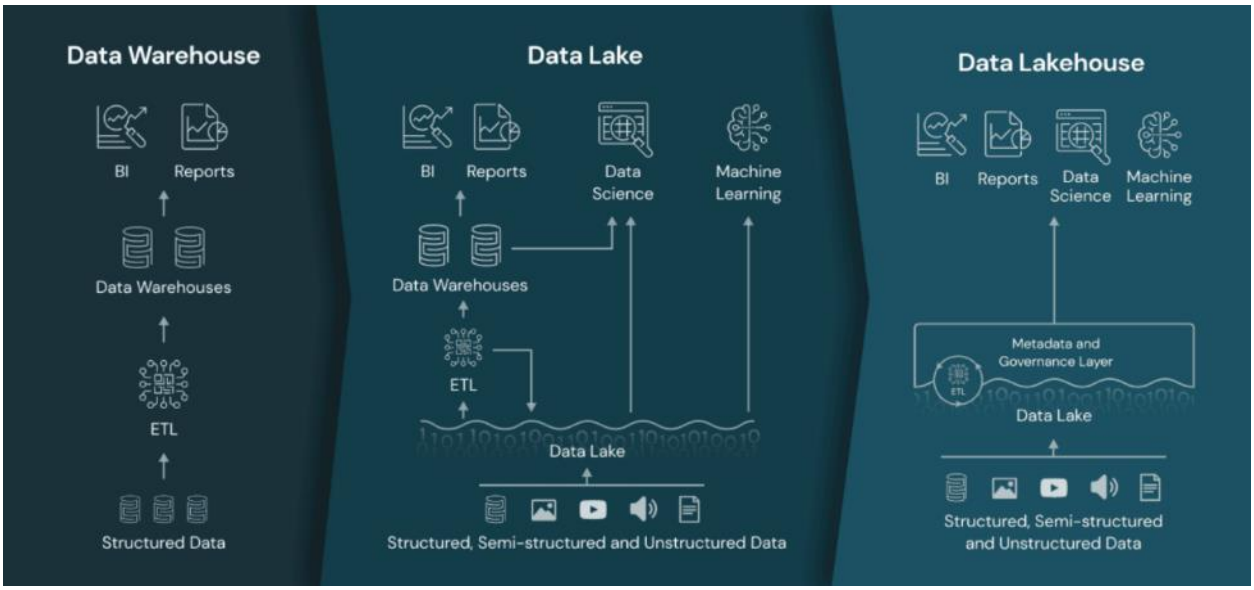

Figure A.1. Comparison of Data Storage Architectures.[1]

A data warehouse is like a well-organized library of information. Everything has a place, and before it is stored, it is cleaned, standardized, and arranged into clear categories that answer common questions consistently, such as addressing key performance indicators or compliance measures. Because data warehouses prioritize structure and quality, they are where decision-makers go for trusted, repeatable answers. A data warehouse is excellent for dashboards, reporting, and repeatable analyses.

A data lake is more akin to a big storage shed for keeping all types of information in its original form. Rather than cleaning and reshaping everything upfront to fit a pre-defined dashboard or analysis, the data lake accepts data as it comes (i.e., sensor feeds, logs, images, documents,

[1] Databricks, "What is a Data Lakehouse?" January 30, 2020, accessed December 1, 2025, https://www.databricks.com/blog/2020/01/30/what-is-a-data-lakehouse.html .

semi-structured records) so analysts can explore and discover value later on an ad hoc basis. Data lakes focus on flexibility and scale, capturing sources quickly without slowing down intake. This makes them well-suited for modern analytics and AI, where teams often need to examine raw information, test new ideas, and build models across many formats.

A data lakehouse blends the strengths of both the library and the storage shed. It keeps flexible, open-format data like a lake but adds organizational rules that enable fast, reliable analysis more like a warehouse. In practice, this translates into an ability to explore raw datasets while also maintaining consistency, transactions, and indexing for accurate results. The data lakehouse helps avoid locking into a single system, supports scalable analytics and AI/ML, and enables teams to produce dependable outputs without giving up the freedom to store varied data types.

In everyday use, these architectures complement each other. Raw information can be ingested into a data lake, trusted metrics can then be curated in a data warehouse, and data lakehouse capabilities can be used for high-performance analytics on open formats. The chosen mix of approaches depends on specific goals and questions that need to be addressed: How quickly are insights needed? How much do the data types vary? What are the reporting and compliance requirements? How is AI expected to use the data?

For Emergency Managers

What to Know:

With data storage architectures, emergency managers can capture raw and real-time feeds (e.g., sensors, GIS, CAD/dispatch) in a data lake, produce trusted, auditable reports in a data warehouse, and run fast analytics and machine learning on open formats with a lakehouse. Employing best data storage practices enables quicker and more complete situational awareness, standardized briefing metrics, and faster after-action analysis. Looking ahead, these foundations support near-real-time forecasting, self-serve data products for partner agencies, and AI assistants that surface risks, optimize resource allocation, and explain recommendations with traceable data lineage.

Opportunities to Advance:

- Determine optimal mix of data storage architectures by taking inventory of all data sources and classifying them by data refresh rate and information sensitivity.
- Establish ingestion pipelines, data quality checks, and disaster recovery procedures.
- Train staff on stewardship and database querying tools.

For Standards and Policymakers

What to Know:

A practical standards model for EM data storage should emphasize openness, portability, and governance to enable data warehouses, lakes, and lakehouses to interoperate safely. Opportunities for achieving this can include:

- Favor open, well-documented ways of storing data so it can move between systems and be safely updated over time.
- Keep a shared data catalog with consistent labels so stakeholders can find what they need and use common standards for geospatial data and cross-agency sharing so systems connect easily.
- Require recording where data came from and who changed or used it, set simple quality rules up front, and follow widely accepted security, privacy, and records-management practices suitable for regulated, multi-agency use.

Opportunities to Advance:

- Promote open, portable data and table formats and require an enterprise catalog with standardized metadata for all new platforms.
- Set baseline interoperability policies (including geospatial APIs, shared schemas, and lineage/audit requirements) that apply across agencies and vendors.
- Establish minimum security, privacy, and records management controls and require vetted cloud services or equivalent assurances.
- Fund cross-agency pilots and shared services to validate streaming ingestion, reproducible analytics, and policy enforcement.
- Create a concise compliance checklist and a certification path to embed these requirements in grants, contracts, and procurements.

For Emergency Management Vendors

What to Know:

Today, most EOCs have operational systems (e.g., CAD/RMS, GIS, mass notification, WebEOC), shared drives and spreadsheets for reporting, and ad hoc data pipelines that cannot keep up with new data sources coming online. Data is often siloed, with inconsistent definitions and limited metadata, making it hard to trust and reuse. Some cloud services may be in place and streaming ingestion is typically minimal.

Products should achieve the following qualities and capabilities:

- Provide ingestion for both streaming and batch data and integrate with common EM systems (e.g., GIS, CAD/RMS, sensors, messaging).
- Expose rich metadata, lineage, and granular access controls.
- Support geospatial and time-series performance, schema evolution, and automated data quality checks.

- Ensure high availability, disaster recovery, and well-documented APIs/SDKs so AI applications can connect predictably and respect governance.

Opportunities to Advance:

- Deliver reference architectures that show lake ingestion, warehouse curation, and lakehouse analytics working together for EM use cases.
- Add first-class support for catalogs and policy enforcement, offer migration tooling from legacy systems, and embed observability and Financial Operations (FinOps) features (i.e., managing cloud costs and aligning with EM objectives).
- Co-develop mission-critical test scenarios with customers, publish use-case libraries and training, and position interoperability with standardized AI integration approaches as a core differentiator.

For AI Leaders

What to Know:

EM blends heterogeneous, high-velocity, geospatial, and sensitive data under strict uptime and audit requirements. Workloads must handle real-time alerts and changing conditions while ensuring reproducibility, provenance, and least-privilege access. Multi-agency coordination, legacy systems, and variable network conditions add complexity. Opportunities include faster risk detection, resource optimization, and explainable recommendations that align with operational playbooks.

Opportunities to Advance:

- Partner with EM organizations to run low-risk pilots first, connecting to both curated warehouse data and raw lake/lakehouse streams via standardized connectors.
- Build reproducible training and inference pipelines with data quality checks, and measure performance with clear objectives (e.g., freshness, latency, availability), accuracy metrics, and human-in-the-loop review.
- Document workflows tailored to EM operations, harden systems for resilience, and scale only after governance, auditability, and operational reliability are consistently met.

A.1.2 Data Access Architectures

Data access architectures (Figure A.2) describe how people and systems discover, share, and control data across an organization and its partners. Two widely utilized approaches are the data fabric and the data mesh. Fabric is primarily a technology layer that connects and governs data across systems, while mesh is an operating model that assigns responsibility for well-defined data products to the domains that know them best.

A data fabric is like a universal map and rulebook. It helps find what exists across data warehouses, data lakes, applications, and the cloud. A data fabric also helps to understand the data's lineage and enforce who can see or use it. Typical capabilities include a catalog to

discover datasets, metadata to describe them, search to locate the right information, and policy controls to make sure data access is safe and compliant. The fabric does not replace existing systems—it coordinates them—so teams can confidently connect AI and analytics to the right data without bypassing governance.

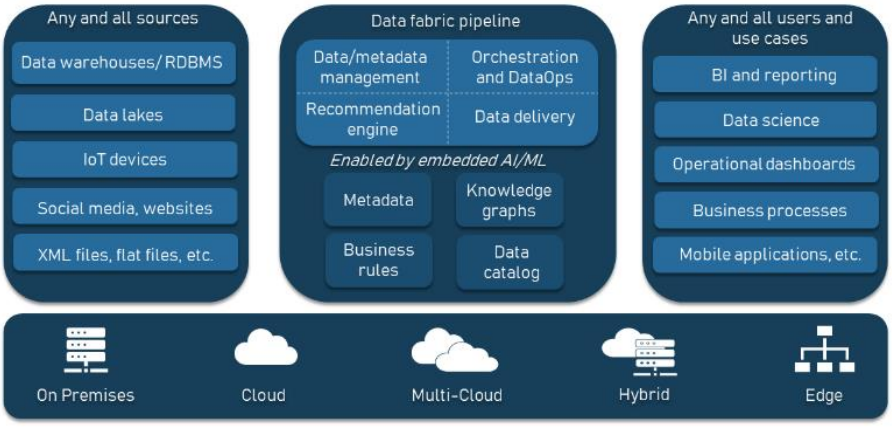

Figure A.2. Data Fabric Architecture Example.[1]

A data mesh shifts ownership closer to the source. Instead of one central team managing all data, each domain (e.g., law enforcement, fire, utilities health) publishes data products with clear definitions, quality expectations, documentation, and access rules (Figure A.3). Shared standards tie these products together so they can be combined across the organization. The mesh recognizes that domain experts understand their data best and are accountable for keeping it accurate and useful. This often speeds delivery, improves relevance, and makes it easier to scale analytics and AI because responsibilities are clear.

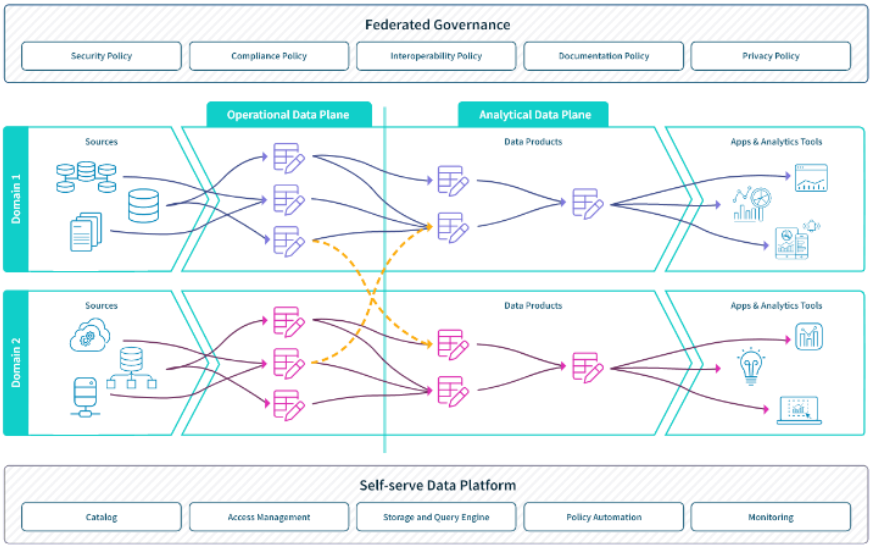

Figure A.3. Data Mesh Architecture Example.[2]

[1] AltexSoft. “Data Fabric Architecture.” August 21, 2022. https://www.altexsoft.com/blog/data-fabric/

[2] Qlik. “Data Mesh Architecture.” 2025. https://www.qlik.com/us/data-management/data-mesh.

Together, fabric and mesh improve both discoverability and accountability. The fabric provides the common catalog, lineage, and policy enforcement that mesh domains use to publish and govern their data products. The mesh ensures that each product has a purpose and an owner on whom others can rely. As a result, data becomes easier to find and safer to use, and teams know who to trust for each dataset. This is crucial for AI adoption, where the right data at the right time used under the right rules determines whether systems are helpful, fair, and reliable.

For Emergency Managers

What to Know:

Data mesh and data fabric architectures provide fundamentally different approaches to organizing and accessing the data sources that EM organizations rely on during emergency operations. Data mesh treats data as a product owned by specific domains (i.e., fire services own fire-related data, EMS owns medical data, public works owns infrastructure data), with each domain responsible for making their data accessible to others while maintaining control and quality. Data fabric creates a unified access layer that intelligently connects all data sources, allowing AI systems and analysts to query across everything without needing to know where data physically resides or how it is formatted. Both approaches solve the problem of data silos that slow emergency response, but mesh emphasizes domain ownership and accountability while fabric emphasizes seamless integration and unified access.

Opportunities to Advance:

- Evaluate whether EM data challenges stem primarily from unclear ownership and accountability (favoring data mesh) or from integration complexity and fragmented access (favoring data fabric).
- Identify domain boundaries within the EM ecosystem and assess whether each domain has the technical capacity to serve their data as products if pursuing a mesh approach.
- Pilot data fabric technology connecting the most critical but fragmented data sources (e.g., CAD, sensors, GIS) to demonstrate unified access benefits before committing to organization-wide implementation.
- Establish data governance policies defining ownership, quality standards, and access permissions that work regardless of whether pursuing mesh (domain-owned) or fabric (centrally integrated) architecture.
- Train EOC staff and IT personnel on the chosen architecture's implications for how they request data access, report data quality issues, and collaborate across organizational boundaries during incidents.

For Standards and Policymakers

What to Know:

Policy frameworks must accommodate both data mesh and data fabric approaches as viable architectures for EM data access, recognizing that different jurisdictions face different organizational structures, technical capabilities, and governance challenges. Data mesh

requires strong domain ownership and accountability, making it suitable for federated EM environments with clear organizational boundaries but requiring policies that ensure domains maintain and share their data products reliably. Data fabric requires robust centralized integration infrastructure and metadata management, making it suitable for consolidated operations but requiring policies that ensure the integration layer does not create single points of failure or compromise data sovereignty. Standards are needed to define what constitutes acceptable data product quality in mesh architectures and what metadata requirements enable effective fabric integration, while liability frameworks must clarify responsibility when data access issues affect emergency response under each architectural approach.

Opportunities to Advance:

- Develop architectural guidance helping EM organizations assess whether data mesh or data fabric better fits their organizational structure, technical maturity, and governance capabilities.
- Establish certification standards for data mesh implementations verifying that domain teams maintain data products meeting quality and availability requirements rather than just claiming ownership.
- Create technical standards for data fabric metadata and integration layers ensuring they do not become proprietary lock-in mechanisms that prevent organizations from switching vendors or adding new data sources.
- Fund pilot programs testing both architectures in multi-jurisdictional EM contexts, documenting which approach better supports mutual aid coordination and regional incident response.
- Develop liability frameworks clarifying responsibility when data access failures affect emergency response, addressing whether domain owners (mesh) or integration platform providers (fabric) bear accountability for outages or quality issues.

For Emergency Management Vendors

What to Know:

Data mesh and data fabric represent different product strategies for solving EM data access challenges, with mesh favoring distributed platforms enabling domain autonomy and fabric favoring centralized integration platforms providing unified access. Vendors pursuing mesh approaches must provide tools enabling non-technical domain teams to easily publish, maintain, and monitor their data products without requiring dedicated data engineering resources. Vendors pursuing fabric approaches must provide robust integration capabilities handling the complex formats, protocols, and quality levels characterizing EM data sources while maintaining acceptable performance during high-volume incident operations. Both approaches require strong governance capabilities, comprehensive metadata management, and clear data lineage tracking, but mesh emphasizes domain-level accountability tools while fabric emphasizes cross-source integration intelligence and unified access interfaces.

Opportunities to Advance:

- Develop data mesh platforms with user-friendly interfaces that enable domain teams to publish data products, monitor usage and quality, and respond to consumer feedback without requiring specialized technical expertise.
- Create data fabric solutions with AI-powered integration capabilities that automatically discover schemas, map relationships, and handle format variations across EM data sources without extensive manual configuration.
- Provide comprehensive governance frameworks supporting either architecture, including data catalogs, lineage tracking, quality monitoring, and access control appropriate for EM security and compliance requirements.
- Establish partnerships with EM organizations to pilot both architectures in real operational contexts, gathering evidence about which approach better supports time-critical incident response and multi-agency coordination.
- Build migration tools and hybrid capabilities, recognizing that organizations may start with one architecture and transition to another or they may need elements of both depending on different data domains and use cases.

For AI Leaders

What to Know:

EM AI systems must access multiple data sources (e.g., structured databases, real-time sensors, geospatial platforms, unstructured documents, external feeds) that are often fragmented, inconsistently formatted, and controlled by different organizational domains. Data mesh architectures require AI systems to discover and consume data products from multiple autonomous domains, handling varying schemas, quality levels, and availability patterns while respecting domain-specific access controls and usage policies. Data fabric architectures provide AI systems with unified access abstractions but require robust metadata understanding and intelligent query federation to efficiently retrieve relevant data without creating performance bottlenecks or violating data sovereignty requirements. Successfully implementing AI in either architecture requires balancing the autonomy and diversity that mesh enables against the integration simplicity that fabric provides, while ensuring AI systems maintain acceptable performance during time-critical emergency operations.

Opportunities to Advance:

- Design AI systems with flexible data access patterns supporting both mesh consumption (discovering and querying autonomous domain data products) and fabric integration (leveraging unified access layers with intelligent metadata navigation).
- Implement comprehensive metadata understanding capabilities enabling AI to interpret schemas, assess data quality, understand lineage, and respect access controls regardless of whether data comes from mesh domains or fabric integration layers.

- Develop efficient query federation and caching strategies ensuring AI systems can retrieve data from multiple mesh domains or fabric sources with acceptable latency during time-critical emergency operations without overwhelming source systems.
- Create transparent data lineage and provenance tracking showing which domains or sources contributed to AI recommendations, supporting accountability requirements regardless of underlying architecture.
- Build graceful degradation capabilities ensuring AI systems continue functioning when individual mesh domains become unavailable or fabric integration layers experience partial failures, adapting recommendations based on available data rather than failing completely.

A.2 Agent Internet

Agent internet establishes the universal connectivity fabric through which agents discover, access, and interact with external data sources, services, and one another. It abstracts transport, authentication, and service-discovery concerns so agents can operate across heterogeneous environments.

A.2.1 Autonomous Systems

Autonomous systems are platforms that can operate independently based on predefined inputs. These platforms can be aerial, vehicular, bipedal, or quadrupedal, with each having different characteristics that can be used by first responders and EM. While the backend and vision interfaces may be proprietary, the textual and voice inputs are typically based on natural language. Integration with other data systems, such as CAD, may allow the autonomous system to be activated without human interaction. Through a world model that defines possible actions, a system that understands the request, and a control system that implements the request, an autonomous system can enact a user's request. The dominant approach uses two systems like the human brain: one to respond to higher-level functions and one to provide more immediate reactions (Figure A.4).

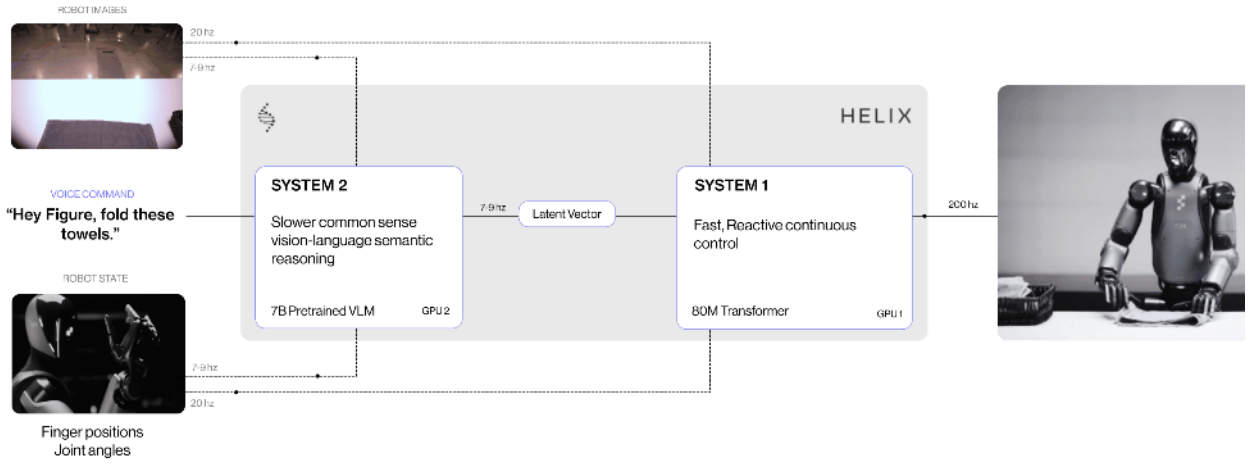

Figure A.4. Two-System Approach by the Figure Helix.[1]

[1] Figure. "Helix." Figure accessed December 1, 2025. https://www.figure.ai/helix

Fundamental challenges exist for truly autonomous systems: they require a precise understanding of their environment, the capabilities of the autonomous system, the training data that allows the autonomous system to understand, the understanding of the request, the ability to break the request into discrete steps, the ability to “remember” the previous decisions, and the ability to interpret the unexpected. All these factors limit the ability for autonomous systems to reliably function in life-safety situations, though there are some exceptions.

Aerial platforms are capable of autonomous operations, with concepts focusing on operations that can be performed more quickly than personnel or in areas that are too hazardous for personnel. In rugged terrain, autonomous drones with thermal imaging or wireless detection technology can be used to locate lost or injured civilians. In HAZMAT situations, autonomous systems could perform the necessary detection without endangering public safety personnel. In disaster situations, autonomous systems can be used to perform damage assessment.

For Emergency Managers

What to Know:

Autonomous systems provide a powerful capability when used within the narrow confines in which they can successfully operate. As time passes, these capabilities will grow and EM organizations should be willing to embrace these technologies. By developing concepts of operation that align with current policies and approaches, performance can be benchmarked. Once the technology is accepted, more leading-edge concepts of operations and policies can be created.

Opportunities to Advance:

- Focus on tasks that can be performed faster and more safely by autonomous systems:
- Determine potential applications that meet the needs of the organization and establish key performance indicators to benchmark autonomous system performance. EM organizations should be open to sharing information that can assist in their mission to autonomous system developers.

For Standards and Policymakers

What to Know:

Autonomous systems require rigorous testing to ensure that they can operate in life-safety situations. Training data for autonomous systems is necessary but may be difficult to make available. Liability is a concern for EM organizations, particularly for systems that do not require human interaction.

Opportunities to Advance:

- Develop public safety datasets and standardized testing methods to help prove reliability in life-safety events. Standards should define methods that allow real-life data from operations to be used without compromising PII and privacy.
- Establish legislation to determine liability in the event an autonomous system acts in an inappropriate manner, so EM organizations have a clear understanding.
- Determine appropriate funding methods and establish new funding as required.

For Emergency Management Vendors

What to Know:

EM vendors must recognize that while a type of autonomous system may be applicable to multiple members of a service, the actual use may vary wildly based on the mission, policies, and geography. AI companies may be unaware of the specific requirements for public safety operations along with the policies that govern them.

Opportunities to Advance:

- Document and communicate in EM the circumstances where their products will operate safely and within the respective limitations.
- Partner with AI leaders as a “technical bridge” to help them understand how public safety operates and help EM organizations better implement novel applications that autonomous systems can provide.

For AI Leaders

What to Know:

Public safety organizations operate in rulesets different from the general public and autonomous systems need to act accordingly. Little training data may be available and scenarios may exist that have never been seen before and require split-second decisions. When an autonomous system meets these types of situations, it needs to gracefully and safely fail in a way that does not increase the danger.

Opportunities to Advance:

- Develop emergency-specific autonomous systems that meet mission-critical reliability standards and security requirements. World models need to be able to learn “on-the-fly” when unexpected situations occur. Developing platform resilience while allowing creative but safe problem-solving opens new opportunities for autonomous systems.

A.2.2 Agent Schema Standards

For AI agents to communicate, they need to understand each other's capabilities. AI agent schema standards provide this interoperability by defining the skills and the inputs and outputs available within each skill (Figure A.5). The expected inputs and outputs are critical to the successful use of AI agents, as there is a risk of misinterpretation between organizations (such as the measurement units for a wind sensor).

Schemas also define the authorization methods available such as OAuth or OpenID Connect, with additional information about the agent's capabilities available after authentication. A standardized authentication method is required to safely interoperate, and some authentication methods may be specific to individual cloud platforms.

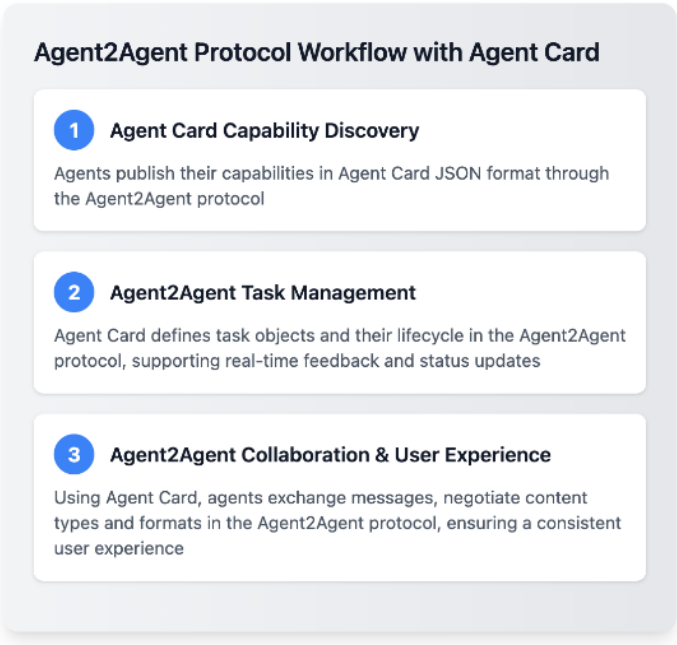

Figure A.5. Agent2Agent Agent Card (Defines the AI Agent's Capabilities).[1]

Agent schema standards may be used to create registries or directories of the available AI agents. Discovery may be built on top of an existing schema (such as the Agent Directory Service using the Open Agent Schema Framework). However, some standards (such as A2A) may not directly support discovery. In these instances, manual or proprietary methods may be required.

Building on the AI agent-to-agent example, wildfire response requires significant coordination between multiple agencies and technologies. Using the built-in AI agent schemas and a directory service, agencies can quickly establish capabilities that pass personnel information and sensor data. For example, an upstream AI agent can request a wind sensor update to decide if a fire crew is in danger. The downstream AI agent can provide the requested wind sensor information as well as detail that the measurement is from thirty minutes ago, allowing the upstream agent to either request an update from the sensor through the downstream agent or to warn the incident commander that the information may be unreliable.

[1] Agent Card. "Features." Figure from Agent Card Platform. Accessed December 1, 2025. https://www.agentcard.net/#features

For Emergency Managers

What to Know:

Agent schema standards allow for the exchange of capabilities to enable interoperability and require significant planning. This involves not only the authentication methods but also the compatibility of transferring information between AI agents. Agent registries/directories allow for easier finding of AI agents, but this may be a manual process. Outputs from the AI agents may not have a timestamp associated with the information on which a decision is being based.

Opportunities to Advance:

- Work together to determine the best agent approaches for a region. Statewide Interoperability Coordinators may be a potential resource for accomplishing this task.
- Wherever possible, agent registries/directories should be automatically created.
- Agent registries/directories should be maintained and kept up to date to avoid failures at critical times.
- If time-dependent information is provided, ensure that the time the measurement was taken is also provided.

For Standards and Policymakers

What to Know:

Having compatible authentication methods is critical to AI agent interoperability. Public Safety Communications (SAFECOM) data approaches may need to be adjusted to accommodate the standardized sharing of AI agent-to-agent capabilities. Public safety agencies expect that one AI agent interface will adapt the information and capabilities based on permissions and the Memorandums of Understanding (MOUs). Having to create duplicates of each interface based on each partner agency is a significant burden that should be avoided.

Opportunities to Advance:

- Ensure that preplans and preconfigured credentials are factored into the agent directories. Federated credentials for authentication are required and should not be specific to an individual platform.
- Schema standards should allow for operations based on MOU policies, allowing for one flexible interface for multiple entities. These standards should clearly indicate what features and the granularity of the information available to each organization.

For Emergency Management Vendors

What to Know:

Connections between agencies may need to be quickly activated or created to respond to a unique situation. EM organizations may not have knowledge of the terms being used by other organizations or where data may be misinterpreted. Test procedures help ensure that the agents have the appropriate credentials and that the agentic workflow is operating correctly.

Opportunities to Advance:

- Help AI leaders in creating the flexibility to rapidly and securely connect external AI agents.
- Leverage relationships with EM organizations to help ensure that deployed AI agent schemas are followed. EM vendors should create test procedures that validate the interoperability and workflow based on realistic configurations.
- Where possible, pipeline views should identify issues within an AI agent chain.

For AI Leaders

What to Know:

The integration of registry/directory capabilities directly into the standards is a critical feature that public safety organizations need. Each public safety organization operates with different information formats and data expectations. Public safety organizations may also have terminology, phrases, or shorthand that may not exist in everyday language. EM vendors may create non-standard configurations specific to their ecosystem.

Opportunities to Advance:

- Work with standards groups to determine the best approach for a standardized solution based on common AI agents.
- Work with public safety organizations to generate template input and output schemas to avoid agent variations for similar tasks.
- Ensure that any natural language approaches understand the public safety vernacular.
- Ensure that the schemas remain EM vendor agnostic.

A.3 Protocols

Protocols define the common languages and interaction patterns that allow agents to communicate, negotiate tasks, share context, and enforce trust. Standardizing message formats, capability declarations, and security primitives ensures interoperability while remaining agnostic to specific implementation stacks.

A.3.1 Agent-to-Agent

An agent-to-agent approach can also reduce the complications in data sharing, such as CAD-to-CAD integration (Figure A.6). While the graphic shows a parallel relationship, it is possible to create a hierarchical relationship as well. Different terminology is used to describe the controlling AI agent, such as orchestrator or supervising agent. A hierarchical approach aligns well with an ICS, allowing a distribution of the tactical versus strategic tasks. By implementing agent-to-agent communications with data centered around standardized types such as NIMS, information can be easily, securely, and quickly shared between agencies as needed.

As an example, wildfire response requires significant coordination between multiple agencies and technologies. The orchestrator AI determines that fire crews are needed at a particular location and Incident Command agrees. The determination of appropriate resources is delegated to a logistics AI that queries the responding agencies for available personnel and equipment. The logistics AI can then interface with a routing AI to determine if the crew can reach the location and estimated time of arrival for each resource.

Tactical decisions can be organized as well. In a wildfire response, the orchestration of fire maps requires the input of multiple interfaces to create an accurate picture of the situation. A fire map AI agent can coordinate with other AI agents that can provide satellite imagery, wind sensor analysis, fire modeling, personnel location, and calculated egress routes to ensure that fire crews operate safely and are not overrun.

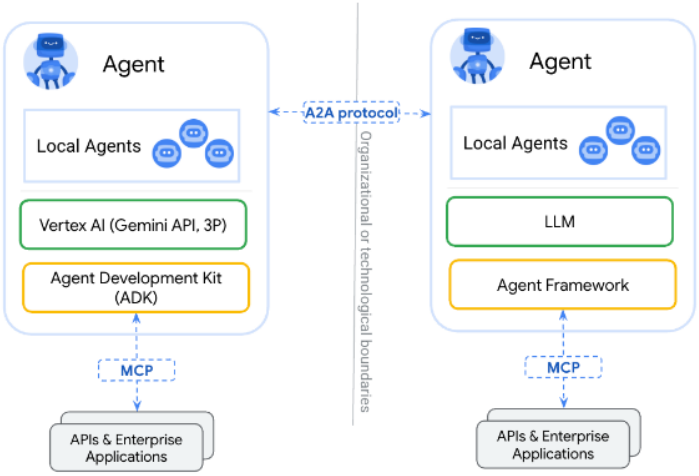

Figure A.6. Agent-to-agent Approach.[1]

[1] "A2A and MCP: Seamless Communication for AI Agents," *Medium*, accessed December 1, 2025, https://arshren.medium.com/a2a-and-mcp-seamless-communication-for-ai-agents-32ca63032659 .

For Emergency Managers

What to Know:

The goal of AI agent-to-agent communications is not to remove the human from the loop but to speed tasks and improve coordination. Opportunities exist not just between similar agencies but also between disparate organizations that interoperate or need to share information.

Opportunities to Advance:

- Focus on tasks that can be done faster and safer by AI agents.
- Identify information flows that routinely have bottlenecks and determine if AI agent-to-agent communications are appropriate.
- Determine appropriate points to allow for human decisions and/or confirmation in the AI agent-to-agent environment.

For Standards and Policymakers

What to Know:

MOUs and Mutual Aid should account for this new approach to information sharing. ICS and SAFECOM communication approaches may need to be adjusted to accommodate AI agent-to-agent communications.

Opportunities to Advance:

- Evaluate current interoperability approaches to determine what changes are required to incorporate the new dataflow. System ownership and guidance should be established to avoid interagency conflicts.
- Determine appropriate funding methods that are available across entities and establish new funding as required.
- Establish legislation to determine liability in the event an AI agent acts in an inappropriate manner, so EM organizations have a clear understanding and expectation.

For Emergency Management Vendors

What to Know:

Integration opportunities are available to provide a turnkey AI agent-to-agent solution, and EM vendors can act as the technical and political bridge between organizations. AI companies

may be unaware of the specific requirements for public safety operations along with the policies that govern them.

Opportunities to Advance:

- Ensure integration ownership of the AI agent-to-agent system must be undertaken to be successful.
- Create an overall AI agent-to-agent system dashboard can help EM organizations know the status of the system and where potential issues may occur.
- Partner with AI leaders as a “technical bridge” to help them understand how public safety operates and help EM organizations better implement novel applications that autonomous systems can provide.

For AI Leaders

What to Know:

Public safety organizations operate in defined methods that also have specific terminology. Understanding the terminology, the outputs of each AI agent, the capabilities of each equipment, and the skills of each person is vital to mission success. When an AI-to-AI agent meets time-critical or life-safety situations, it needs to gracefully and safely fail in a way that does not increase the danger.

Opportunities to Advance:

- Develop emergency-specific AI agent-to-agent systems that meet mission-critical reliability standards and security requirements. Public safety operations center around people, and AI agent-to-agent systems should be designed not to replace personnel but to reduce their burdens. AI agent-to-agent system designs should include humans in the loop for critical decisions.

A.3.2 Agent Gateways

Agent gateways allow for secure communications between AI agents and other data protocols, such as MCP and OpenAPI (Figure A.7). A key component of the agent gateway is the protection of the AI agents and the data sources. Standard approaches include using authentication and Role-Based Access Control to restrict access. Gateways also protect against sophisticated attacks that can compromise the agent system or leak data such as prompt injections and tool poisoning attacks.

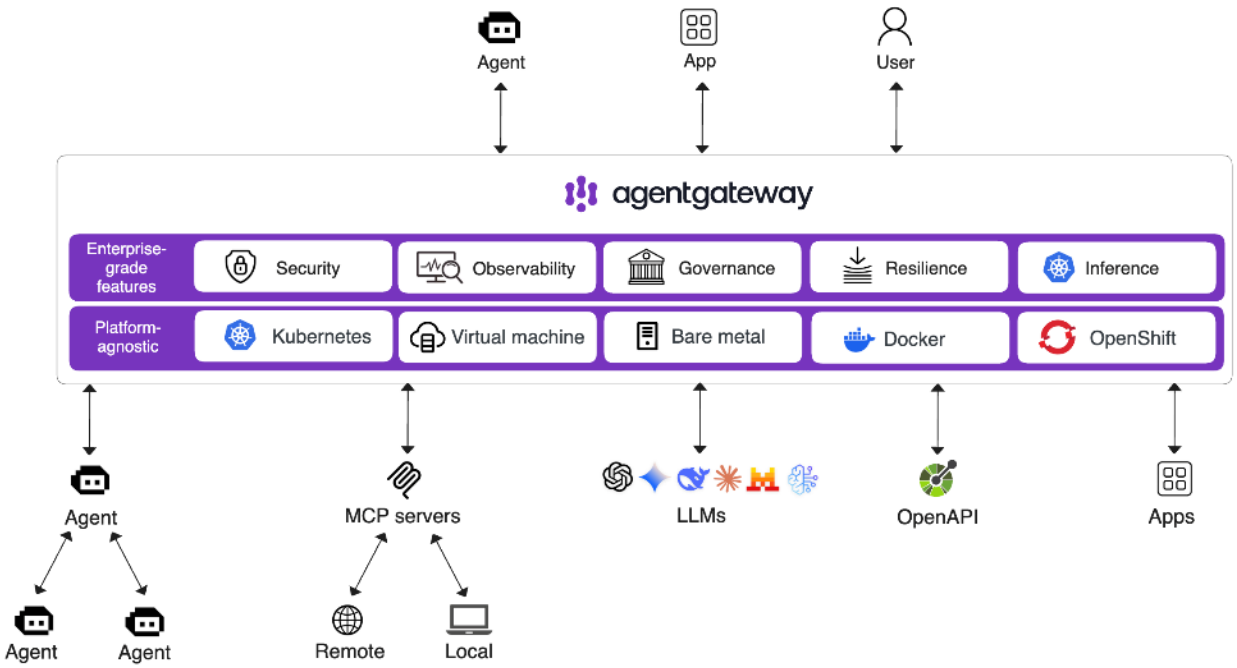

Figure A.7. AI Agent Gateway.[1]

Additionally, agent gateways can assist in resiliency and scaling by using dynamic routing for requests, which is critical for failure situations such as natural disasters or for changing situations such as an expanding wildfire. Dynamic routing can also allow for graceful fallback should an AI agent or other data source have an issue.

As response requires significant coordination between multiple agencies and technologies, secure connections may need to be rapidly established between agencies or merely "turned on" if existing MOUs are in place, with agent gateways performing authentication. Data sources may need to be rapidly integrated from outside the organization and potentially untrusted data sources, such as wind sensors and agent gateways, would help minimize the risks of these data sources.

For Emergency Managers

What to Know:

In the event of a failure or delay with a data source or another AI agent, the AI agent may make decisions with incomplete or stale information. The agent gateway dashboard is critical to understanding the status of the overall AI agent system.

Opportunities to Advance:

- Design dashboards to quickly relay information for personnel verification and fallback procedures.

[1] Agent Gateway, "Architecture," accessed December 1, 2025, https://agentgateway.dev/docs/about/architecture/ .

- Establish policies for handling failure or information delay scenarios.
- Train personnel to activate new data sources based on established policies.
- Regularly practice procedures to ensure readiness and reliability.

For Standards and Policymakers

What to Know:

The AI agent gateway is a critical component to the security of the AI agent system, and AI agents should not be directly exposed. MOUs and Mutual Aid should account for this new approach to information sharing.

Opportunities to Advance:

- Ensure all requests for proposals for AI agents include protections such as agent gateways.
- Create preplans and preconfigured credentials to enable rapid integration of external data sources and AI agents.
- Collaborate with EM organizations to establish appropriate responses for delayed or stale time-critical information.
- Define policies for excluding outdated data from AI agents' decisions (e.g., when specific data types exceed a time threshold).

For Emergency Management Vendors

What to Know:

AI agent gateways should be an integral part of a turnkey AI agent-to-agent solution. Connections between agencies may need to be quickly activated or created to respond to a unique situation. Data sources may require access to the internet or with untrusted data partners and appropriate safeguards are required.

Opportunities to Advance:

- Create an overarching AI agent-to-agent system dashboard to monitor system status and identify potential issues.
- Design mechanisms for rapid integration of external data sources.
- Collaborate with EM vendors to evaluate the risk versus reward of integrating data partners into the AI agent framework.

For AI Leaders

What to Know:

Public safety organizations will be targeted and may not be aware of the methods necessary to protect AI agents. AI leaders need to understand the vulnerabilities with not only the AI agents in use but also the data sources that EM organizations use. In dynamic situations, stale data is just as dangerous as missing information.

Opportunities to Advance:

- Proactively address the latest security threats faced by EM organizations.
- Ensure a deep understanding of the information used by AI systems to strengthen security measures.
- Flag stale or delayed requests with time information to ensure AI agents make decisions based on current data.
- Clearly communicate real-time data substitutions (e.g., alternative data sources due to failure) to EM organizations to enhance decision-making.

A.3.3 Model Context Protocol

MCP is a standardized framework developed by Anthropic that revolutionizes how AI applications interact with external data sources and tools (Figure A.8). Instead of requiring custom, time-intensive integrations for each data source, MCP provides a universal interface that enables AI models to seamlessly access databases, file systems, APIs, web services, and enterprise applications through a consistent, secure methodology.

MCP operates through a client-server architecture where "MCP servers" act as specialized bridges between AI applications and data sources. These servers handle critical functions, including authentication, data formatting, security protocols, and access control, ensuring AI systems can only access authorized information within defined parameters. The protocol supports real-time data streaming, tool execution, and dynamic resource management while maintaining strict security boundaries and audit trails.

The framework addresses a fundamental challenge in enterprise AI adoption: integration complexity. Traditional approaches require months of custom development work and ongoing maintenance for each AI-data source connection. MCP eliminates this barrier by standardizing the integration process, allowing organizations to deploy MCP servers that instantly connect existing infrastructure to AI applications.

This standardization promotes unprecedented interoperability between different AI tools and data ecosystems, accelerating AI adoption across industries. Organizations can leverage AI capabilities without rebuilding their technology stack, while maintaining enterprise-grade security, governance, and compliance requirements. MCP essentially democratizes AI integration, making sophisticated AI applications accessible to organizations regardless of their technical resources or development capacity.

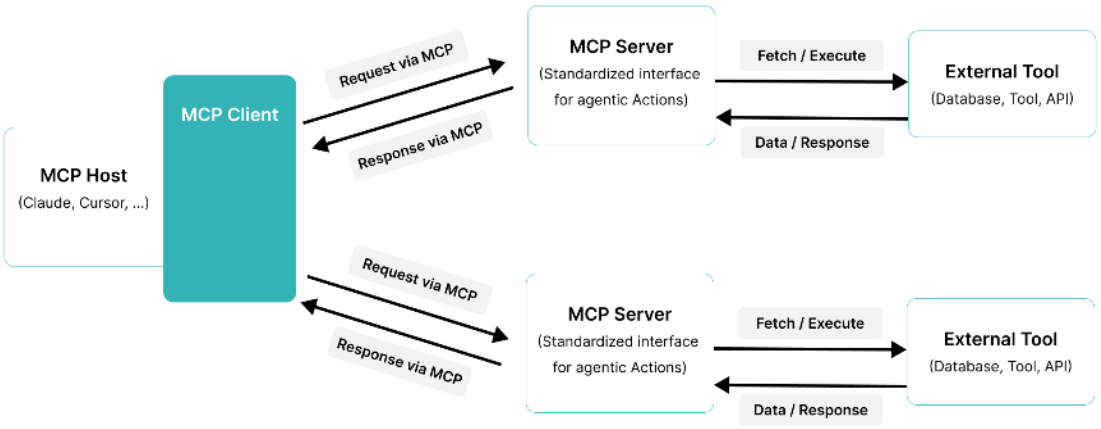

Figure A.8. Model Context Protocol Example.[1]

For Emergency Managers

What to Know:

MCP enables EM organizations to integrate AI with existing systems like databases, sensor networks, GIS platforms, and communication tools without complex custom development. This allows AI to provide real-time situational awareness, predictive analytics for resource allocation, automated threat detection, and coordinated response recommendations. Emergency managers gain enhanced decision-making capabilities by connecting AI to their entire technology ecosystem, improving response times and operational effectiveness during critical incidents.

Opportunities to Advance:

- Conduct initial inventory of their data sources and systems.
- Deploy MCP servers to connect databases, sensors, GIS platforms, and communication tools once inventory is complete.
- Configure appropriate security permissions and access controls, select MCP-compatible AI applications, and establish data governance policies.
- Train staff on AI integration workflows.
- Conduct pilot testing with non-critical scenarios before full deployment during actual emergency operations.

[1] Dida, "A Practical Introduction to the Model Context Protocol (MCP)," accessed December 1, 2025, https://dida.do/blog/a-practical-introduction-to-the-model-context-protocol-mcp .

For Standards and Policymakers

What to Know:

MCP requires comprehensive policy frameworks addressing data governance, privacy protection, and security standards for AI-system integration across organizations. Standards organizations must develop interoperability protocols, certification processes, and compliance guidelines. Policymakers need regulatory oversight mechanisms for AI data access, liability frameworks for automated decisions, and cross-agency coordination standards. Public sector adoption requires transparency requirements, audit trails, and accountability measures to ensure responsible AI deployment while maintaining public trust and operational integrity.

Opportunities to Advance:

- Support MCP pilot programs in EM.
- Establish interoperability requirements for government systems.
- Create regulatory incentives for adoption.
- Collaborate with standards organization on best practices, security guidelines, and training standards while promoting public-private partnerships that accelerate MCP integration across EM ecosystems nationwide. Standards organizations can develop emergency-specific MCP protocols, certification programs for vendors, and data sharing frameworks for multi-agency coordination.

For Emergency Management Vendors

What to Know:

EM vendors must recognize that MCP compatibility will likely become a competitive requirement as customers demand AI integration capabilities. Early adoption provides market advantage and partnership opportunities with AI companies. Vendors need to understand MCP's technical architecture, develop server implementations for their platforms, and prepare for shifting customer expectations toward AI-enabled solutions. Those who delay risk losing market share to competitors offering seamless AI integration through MCP-compatible EM systems.

Opportunities to Advance:

- Prioritize developing robust MCP server implementations for their platforms, updating software architectures to support AI integration, and creating comprehensive APIs for data access.
- Establish partnerships with AI providers, invest in developer documentation and SDKs, and provide customer training programs.
- Conduct extensive testing with EM clients, develop use-case libraries, and position MCP compatibility as a key differentiator in their marketing and sales strategies.

For AI Leaders

What to Know:

EM organizations operate high-stakes, life-critical environments requiring ultra-reliable data access from different sources, including sensors, GIS platforms, and communication networks. They have complex legacy infrastructures, strict security requirements, and need real-time interagency coordination. AI leaders must ensure MCP implementations meet stringent reliability standards, support mission-critical uptime, handle sensitive data appropriately, and integrate seamlessly with existing EM workflows without introducing operational vulnerabilities or delays.

Opportunities to Advance:

- Develop emergency-specific MCP server implementations that meet mission-critical reliability standards and security requirements.
- Establish partnerships with EM organizations for real-world testing and feedback, create specialized training programs and documentation tailored to EM workflows, and build dedicated support channels.
- Invest in understanding EM data requirements and develop EM-focused use cases.
- Ensure their frontier models can handle the unique demands of life-critical decision-making scenarios.

A.3.4 Data Interoperability Standards

EM relies on rapid information sharing across multiple agencies, jurisdictions, and technology systems. During a major incident, data must flow seamlessly between services monitoring outdoor conditions, 9-1-1 centers, hospitals, utility companies, federal agencies, and countless other partners. Data interoperability standards establish common formats and models that allow different systems to exchange information without requiring custom integration for every connection (see Figure A.9). These standards function as a universal language that lets different computer systems understand each other, much like how standardized shipping containers revolutionized global trade by working with any truck, train, or ship.

Without these standards, organizations face vendor lock-in, where they become dependent on a single technology provider, and data silos, where critical information remains trapped in incompatible systems. During Hurricane Katrina, incompatible communication systems hampered coordination between agencies. Modern interoperability standards prevent these failures by ensuring that incident data, resource requests, and situational awareness information can be shared instantly across organizational boundaries. Open table formats like Parquet and Delta allow data to be stored and accessed by any tool, while domain-specific standards like the NIEM provide pre-built vocabularies for EM concepts.

For AI systems to be effective in EM, they must access and integrate data from multiple sources in real-time. Canonical data models and event schemas ensure that AI agents can correctly interpret information from different systems (e.g., understanding that "incident" means the same

thing whether it comes from a fire department, police dispatch, or EOC). These standards make it possible for AI to analyze patterns across jurisdictions, coordinate multi-agency responses, and provide decision support without requiring emergency managers to manually translate data between incompatible formats.

Figure A.9. Interoperability Continuum.[1]

For Emergency Managers

What to Know:

Data interoperability standards enable agencies to share information rapidly with mutual aid partners, state and federal agencies, and private sector organizations during incidents. Standards like NIEM are specifically designed for EM and public safety, providing ready-made data structures for common concepts like incidents, resources, and facilities. Adopting open formats prevents situations where a user cannot access a given organization's own data because a vendor went out of business or changed their pricing. These standards are critical for multi-jurisdictional incidents where dozens of agencies must coordinate using different technology systems, and they form the foundation that allows AI systems to integrate data from multiple sources to support decision-making.

Opportunities to Advance:

- Inventory current data systems and identify which ones use proprietary formats that create barriers to information sharing with partner agencies.
- Adopt NIEM-compliant data schemas for incident reporting, resource management, and situational awareness to ensure compatibility with state and federal systems.

[1] Voss, B. & Anderson, E. (2019). Interoperability of Real-Time Public Safety Data: Challenges and Possible Future States. https://doi.org/10.6028/NIST.IR.8255

- Establish data exchange agreements with mutual aid partners based on standardized formats, particularly for automatic sharing of resource availability and incident status.
- Include interoperability requirements in procurement specifications, requiring vendors to support open data formats and industry standards rather than proprietary formats.
- Participate in regional data sharing pilots that test standardized information exchange during exercises and small-scale incidents before relying on them during major events.

For Standards and Policymakers

What to Know:

Mandating data interoperability standards through policy and procurement requirements accelerates adoption across EM organizations and levels the playing field for technology vendors. Standards like NIEM already exist for EM but require policy-backing to drive implementation. Open format requirements prevent public agencies from becoming locked into proprietary systems that create long-term costs and operational dependencies. Interoperability standards also enable AI innovation by ensuring that new tools can access existing data without requiring agencies to replace their entire technology infrastructure. Policy frameworks should balance standardization with flexibility, allowing agencies to adopt new technologies while maintaining the ability to share information across organizational boundaries.

Opportunities to Advance:

- Develop procurement language requiring that EM systems support open data formats and NIEM standards as a condition for government contracts.
- Create certification programs that verify vendor compliance with interoperability standards, making it easier for agencies to identify compatible solutions.
- Establish grant incentives that prioritize funding for agencies implementing standards-based data sharing with regional partners.
- Mandate that federally funded EM systems publish their data schemas and APIs to enable integration with state and local systems.
- Convene working groups bringing together EM practitioners, technology vendors, and standards organizations to update data models as new technologies and threat scenarios emerge.

For Emergency Management Vendors

What to Know:

Supporting data interoperability standards expands the potential market by making solutions compatible with the technology ecosystems that EM agencies operate. Agencies increasingly view interoperability as a requirement rather than a feature, and standards compliance differentiates vendors who can integrate into complex multi-agency environments from those who create isolated silos. NIEM and other domain standards reduce development costs by

providing pre-built data models rather than requiring the design of proprietary schemas. Open format support also future-proofs solutions, ensuring they remain valuable as agencies adopt AI and analytics tools that need access to historical data. Vendors who lead in interoperability create network effects where their solutions become more valuable as adoption grows.

Opportunities to Advance:

- Implement support for NIEM data standards in core products, particularly for incident management, resource tracking, and situational awareness capabilities.
- Provide well-documented APIs using industry-standard protocols like REST and GraphQL[1] to enable integration with AI systems and third-party tools.
- Offer data export and import functionality supporting open table formats like Parquet and Delta[2] to prevent customer lock-in and enable analytics.
- Participate in interoperability testing events and certification programs to demonstrate compliance and build trust with government customers.
- Develop partnerships with complementary vendors to create pre-integrated solutions that showcase the value of standards-based interoperability for EM workflows.

For AI Leaders

What to Know:

AI systems for EM must integrate data from highly heterogeneous sources, including legacy systems, modern cloud platforms, Internet-of-Things sensors, and external data feeds. EM organizations have limited IT resources and cannot custom-integrate AI tools with every data source, making standards support essential for practical deployment. Domain standards like NIEM provide semantic understanding of EM concepts, allowing AI systems to correctly interpret incident types, resource categories, and organizational relationships without extensive customization. Open table formats enable efficient batch processing of historical data for model training, while canonical schemas ensure consistent feature engineering. Supporting these standards is not only a technical requirement but also a prerequisite for AI adoption in an ecosystem where interoperability and multi-agency coordination are fundamental to operations.

Opportunities to Advance:

- Design AI systems with native support for NIEM data structures, enabling out-of-the-box integration with EM platforms used by state and federal agencies.

[1] Representational State Transfer (REST) and GraphQL are two popular approaches for designing APIs that facilitate communication between applications. For more information, see https://aws.amazon.com/compare/the-difference-between-graphql-and-rest/.

[2] Parquet is a storage format design for efficiency. Delta Lake builds on Parquet with additional features that can be more robust for data integrity and management. For more information, see https://delta.io/blog/delta-lake-vs-parquet-comparison/.

- Implement adapters for common open table formats (e.g., Parquet, Delta, Iceberg) to efficiently process the large historical datasets emergency agencies maintain for trend analysis and model training.
- Develop schema mapping tools that automatically translate between proprietary formats and standardized models, reducing the integration burden on EM organizations.
- Publish comprehensive documentation showing how AI systems consume and produce standardized data, including examples using real EM scenarios and NIEM vocabularies.
- Participate in EM standards working groups to ensure evolving AI capabilities are reflected in data models and to understand domain-specific requirements that affect AI system design.

A.3.5 Data Provenance and Trust Frameworks

In EM, decisions have life-or-death consequences, and decision-makers must be able to trust and verify the information they receive. When an AI system recommends an evacuation route, predicts flood levels, or suggests resource allocations, emergency managers need to know where that information came from, how it was processed, and whether it can be relied upon. Data provenance and trust frameworks provide standardized methods to track the complete history of data and AI-generated outputs (see Figure A.10), from original sensors and sources through all transformations and model inferences to final recommendations. Provenance is like a chain of custody for information, similar to how evidence is tracked in legal proceedings, ensuring that every piece of data can be traced back to its origin and every decision can be explained and audited.

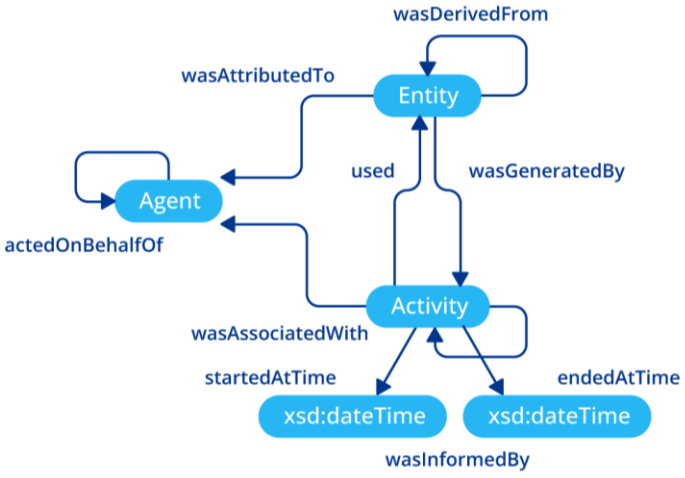

Figure A.10. Data Provenance Example.[1]

The stakes are particularly high when AI systems integrate information from multiple sources or generate synthetic content. Without provenance tracking, it becomes impossible to verify whether a warning came from an official or an unreliable source, whether a damage assessment photograph was taken at the incident scene or was AI-generated, or whether a resource request was authorized by a legitimate incident commander. Standards like C2PA

[1] Astera. "How Does Data Provenance Work." March 3, 2025. https://www.astera.com/type/blog/data-provenance/

provide cryptographic methods to watermark and authenticate digital content, while tracking the components and training data that went into building AI models themselves. These standards create an audit trail that enables accountability and builds trust in AI-assisted decision-making.

Provenance frameworks become especially critical during after-action reviews and legal proceedings following major incidents. Emergency managers must be able to reconstruct the information flow that led to critical decisions, demonstrating that they acted on the best available information and followed proper procedures. When AI systems are involved, this means documenting which models were used, what data they accessed, how confident their predictions were, and what sources informed their recommendations. Strong provenance practices also protect against data quality errors and adversarial attacks, allowing emergency managers to quickly identify and isolate compromised or unreliable information sources before they affect operations.

For Emergency Managers

What to Know:

Data provenance and trust frameworks enable verification of the reliability of AI-generated recommendations before acting on them during high-stakes incidents, providing the audit trails needed to defend decisions during after-action reviews and legal proceedings. These frameworks track where every piece of information originated, how AI models processed it, and what sources informed each output, allowing users to distinguish between official data sources and unreliable information. Content provenance standards like C2PA can verify that situation photographs and videos are authentic rather than manipulated or AI-generated, while model provenance standards document which AI systems produced which recommendations and what training data shaped their capabilities. This traceability is essential for maintaining accountability when AI assists with life-safety decisions and for building public trust in technology-supported EM operations.

Opportunities to advance:

- Implement provenance tracking requirements for all AI systems integrated into an EOC, ensuring every AI-generated recommendation includes source citations and confidence levels.
- Establish verification protocols requiring staff to check provenance metadata before acting on AI recommendations during critical incidents, particularly for evacuation decisions and resource deployments.
- Maintain audit trails documenting which AI models were active during each incident, what data sources they accessed, and what outputs they produced to support after-action reviews and legal accountability.
- Train EOC staff to interpret provenance metadata and recognize warning signs of unreliable or unverified information in AI-generated outputs.
- Conduct tabletop exercises that include scenarios where AI systems receive compromised data, testing an organization's ability to identify and isolate unreliable information sources using provenance tracking.

For Standards and Policymakers

What to Know:

Regulatory frameworks for AI in EM must mandate provenance and traceability to ensure accountability and public trust in life-safety decisions supported by AI systems. Standards like C2PA for content authenticity and SBOM for AI model transparency provide technical foundations but require policy backing to drive adoption across government systems and vendor ecosystems. Provenance requirements enable legal liability frameworks by documenting the information chain that led to decisions, protecting both emergency managers who followed proper protocols and holding accountable those who deployed unverified AI systems. Policy must balance transparency requirements with operational security, ensuring that provenance systems reveal enough information for accountability without exposing vulnerabilities that adversaries could exploit. These frameworks also support public records requirements and freedom of information obligations while protecting sensitive operational details.

Opportunities to Advance:

- Mandate that AI systems procured for government EM include standardized provenance capabilities compliant with C2PA and SBOM frameworks.
- Establish liability protections for emergency managers who follow documented decision protocols using AI systems with proper provenance tracking, while creating accountability for agencies that deploy unverified AI tools.
- Develop certification programs that verify AI vendor compliance with provenance standards, including requirements for model cards documenting training data sources and known limitations.
- Create regulatory requirements for AI-generated public communications to include provenance watermarking, enabling citizens to verify the authenticity of emergency alerts and official information.
- Fund pilot programs testing provenance frameworks in multi-agency emergency operations, documenting best practices for balancing transparency with operational security requirements.

For Emergency Management Vendors

What to Know:

Building provenance and trust capabilities into solutions differentiates products in a market where government customers increasingly demand accountability and auditability for AI-assisted decisions. EM agencies face legal liability for decisions made during incidents, and they require systems that can document the information chain supporting those decisions to protect against litigation. Implementing standards like C2PA for content authentication and maintaining detailed logs of AI model inputs and outputs demonstrates commitment to responsible AI deployment. Provenance features also protect reputation by enabling rapid

identification of data quality issues or model failures, allowing vendors to isolate problems before they cascade across customer deployments. Vendors who lead in transparency and traceability build trust with risk-averse government customers and position themselves favorably as regulatory requirements for AI accountability mature.

Opportunities to Advance:

- Implement C2PA content provenance in any systems that generate or process images, videos, or documents, enabling emergency managers to verify the authenticity of situation reports and damage assessments.
- Maintain comprehensive logs documenting which AI models produced which outputs, what data sources they accessed, and confidence levels for each prediction, with retention policies meeting government audit requirements.
- Publish detailed model cards and SBOMs for all AI systems, documenting training data sources, known limitations, and performance characteristics in EM scenarios.
- Develop user interfaces that prominently display provenance information alongside AI recommendations, making it easy for operators to verify sources and assess reliability during high-pressure incidents.
- Establish partnerships with standards organizations and participate in certification programs demonstrating provenance capabilities meet government accountability requirements.

For AI Leaders

What to Know:

EM represents a high-stakes domain where AI systems must be fully explainable and auditable, requiring provenance tracking throughout the entire AI life cycle from training data collection through model inference to output generation. Decision-makers need to understand not only what AI recommends but also why it made that recommendation, what data informed it, what assumptions were embedded in training, and how confident the model is in its output. Implementing standards like Software Package Data Exchange (SPDX) and SBOM for AI models documents the components and data sources that shaped a system's capabilities, while C2PA integration ensures that AI-generated content can be authenticated and distinguished from human-created or unverified information. Provenance capabilities are not optional features but fundamental requirements for responsible AI deployment in domains where lives depend on decision quality, and building these capabilities from the ground up is far easier than retrofitting them into existing systems.

Opportunities to Advance:

- Instrument AI systems to automatically generate and attach provenance metadata to all outputs, including source citations, model versions, confidence scores, and processing timestamps in standardized formats.

- Implement C2PA signing for any AI-generated content, including text reports, images, and synthetic data, enabling emergency managers to verify authenticity and distinguish AI outputs from human-created content.
- Publish comprehensive SBOMs documenting all training data sources, third-party components, and model architectures, with regular updates as models are retrained or fine-tuned for EM applications.
- Design AI systems to maintain immutable audit logs tracking every data access, model inference, and output generation with cryptographic integrity protections meeting government accountability standards.
- Develop explainability interfaces that surface provenance information in actionable formats for emergency managers, showing decision trees, confidence intervals, and alternative scenarios rather than opaque model outputs.

A.4 Tooling and Enrichment

Tooling and enrichment supplies modular extensions (e.g., function call wrappers, data transformers, and execution sandboxes) that expand or refine an agent's native abilities. This layer enables rapid incorporation of new skills and resources without altering core reasoning components.

A.4.1 Data Discovery for Improved Accuracy: Fine-Tuning

Foundational model LLMs develop their general language skills with pre-training (Figure A.11). In this phase, the LLM learns using neural networks and self-supervised learning methods. Data for this phase is massive in terms of volume and topics. Pre-training data comes from books, journal articles, news articles, websites, and online discourse, among other sources.

However, these more general models can suffer from concerning drawbacks such as hallucinations and missed context. To overcome these drawbacks, the models can be fine-tuned with custom, domain-specific data. This process currently takes days to weeks as opposed to weeks to months for pre-training.[1] Benefits of fine-tuning include increased relevance and accuracy and decreased hallucinations. Additionally, fine-tuning has the capacity to improve data privacy, computational infrastructure security, user experience, model behavior, scalability, and flexibility.

Five types of fine-tuning methods exist: supervised fine-tuning, instruction-tuning, transfer learning, few-shot learning, and reinforcement learning from human feedback. Each method varies in its use of labeled data and how that data receives its labels. Instruction-tuning in particular can be useful when an organization has a list of specific tasks that will be repeated across time and users. Instruction-tuning enables the LLM to learn how to follow specific instructions from a human user and better understand the outputs requested by a specific query.[2]

[1] Parthasarathy, V., Zafar, A., Khan, A., & Shahid, A. (2024). (tech.). *The Ultimate Guide to Fine-Tuning LLMs from Basics to Breakthroughs: An Exhaustive Review of Technologies, Research, Best Practices, Applied Research Challenges and Opportunities* (pp. 6–13). Dublin, Ireland

[2] Nieto, A. (2025, August 7). *LLM pre-training and Custom LLMs*. Databricks. https://www.databricks.com/blog/llm-pre-training-and-custom-llms

As the data for most fine-tuning is domain-specific,[1] they are usually provided by the end user's organization or similar sources. These may take the form of an organization's reports, notes, and emails and include relevant data from other collaborating entities as well. Still, pre-processing steps, including data cleaning, bias adjudication, privacy, and security, are important to undertake. These steps can be cooperative between the end-user organization and the third party performing the fine-tuning, aka the LLM developer. They may also be pre-built into LLM data discovery pipelines, although it is important to still apply domain-specific knowledge as needed to interpret how nuances in data may need to be addressed before the data can be used.

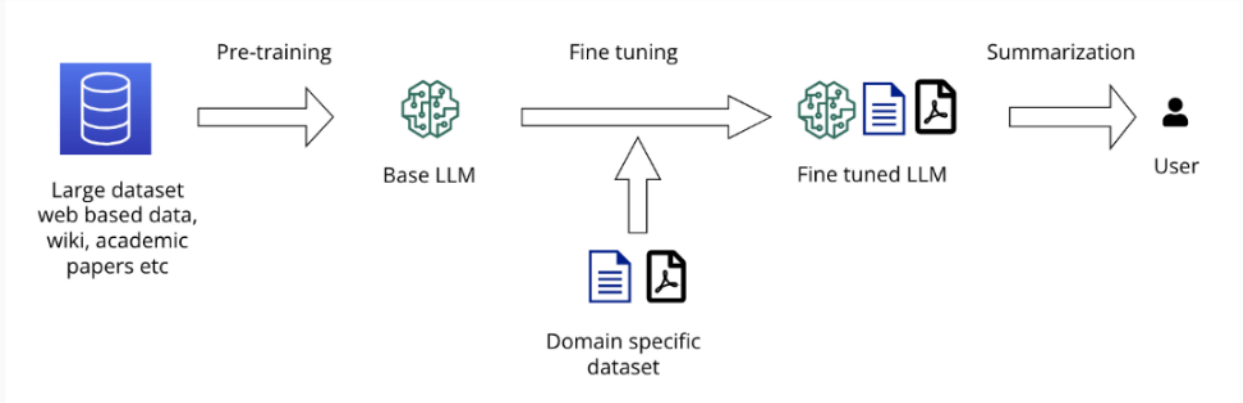

Figure A.11. LLM Training and Fine-Tuning Sequence.[2,3]

For Emergency Managers

What to Know:

While a foundational model may be capable of conducting general tasks such as drafting emails and summarizing notes, EM organizations should consider using a fine-tuned model for tasks such as risk analysis and evacuation planning. In the high-consequence environment in which EM operates, these enhanced models will decrease the risks of producing inaccurate or hallucinated LLM outputs. Even for the more general tasks, a fine-tuned model may mitigate data privacy and security concerns prevalent in the EM realm.

Opportunities to Advance:

- Determine the scope of the LLM they would like to deploy (what questions and tasks should it be able to answer) as the primary step.

[1] While data and prompts can be tailored to a specific domain for instruction-tuning, this method of fine-tuning technically encompasses any attempt to teach a model to follow any set of instructions.

[2] Darga, A., Holkar, P., & Pratap, V. (2025, June 12). *Preparing data for fine-tuning LLMs for contract analysis using Data Prep Kit (DPK)*. IBM developer. https://developer.ibm.com/learningpaths/get-started-data-prep-kit/dpk-llm-applications/dpk-prepare-data-fine-tuning-llms/

[3] Reinforcement learning from human feedback, unlike other types of fine-tuning, does not adhere to this workflow as it occurs once the user is involved.

- Audit internal databases and other sources of data (notes, emails, etc.) to identify relevant data for the in-scope queries following the determination of scope
- Identify collaborating organization and agency data that may additionally enhance the LLM's response quality by enabling it to learn patterns related to specific entities, geographies, processes, etc.
- Work closely with the LLM developer to enable easier cataloging, labeling, and preprocessing of data while also ensuring the incoming data meets required quality standards.

For Standards and Policymakers

What to Know:

Fine-tuning techniques require access to data specific to organizations and agencies. The fine-tuning process becomes more powerful as it is exposed to a higher quantity of data across entities, as it can then learn patterns related to differences (e.g., geographic, policy-driven, demographic) between those entities. Formal agreements can help data sharing efforts adhere to entity and governmental (local, state, and federal) policies while mitigating privacy and security concerns. Techniques like anonymization and encryption can enhance privacy and security as well. Policymakers should also be aware that fine-tuning can propagate dataset biases and other ethical concerns when data of poor quality are used.

Opportunities to Advance:

- Establish data-sharing agreements that clearly outline permissions, scope, and privacy will be necessary.
- Establish workshop groups and conduct workshops to fully discover relevant data sources.
- Establish interoperable data formats and standardized data-sharing protocols to facilitate easier data sharing.
- Consider creating ethical and quality guidelines for those creating, maintaining, and sharing EM datasets.
- Consider data sharing practices that enable entities to maintain ownership and control of their own data while enabling visibility to other entities.

For Emergency Management Vendors

What to Know:

Vendors' ability to help EM clients discover and utilize domain-specific data for model fine-tuning will significantly impact tool accuracy, relevancy, and reliability for EM tasks. Vendors can help by identifying relevant open-source datasets in addition to assisting EM client organizations with leveraging their own data. Establishing partnerships with other entities may provide benefits to EM clientele by adding access to relevant but closed-source data. Vendors

will need to check dataset quality in addition to helping clients with secure data pipelining into vendor tools. Vendors may find it helpful to make recommendations about dataset interoperability and standardization and assist with implementation. Due to the continuous evolution of the EM domain, client feedback and iterative fine-tuning processes will likely be most beneficial.

Opportunities to Advance:

- Work closely with EM organizations to enable them to more easily digitize, catalog, label, and preprocess data as needed while also ensuring the incoming data meets required quality standards.
- Identify and recommend strategies for standardizing data collection in interoperable data formats.
- Fine-tune models for EM applications with open-source incident and other relevant data.
- Consider establishing partnerships with third parties (local governments, nongovernmental organizations, etc.) to enhance data availability for building and tuning EM-specific tools.
- Consider iterative fine-tuning and feedback gathering.

For AI Leaders

What to Know:

Mission-critical EM tasks will necessitate generative AI models beyond foundational capabilities, so advancing methods for incorporating domain-specific knowledge into model outputs will be imperative. Data for EM tasks comes not only from the end user's organization but also from academia, government, private sector, non-governmental organizations, and crowdsourcing. As trust, privacy, and security are of utmost importance in EM, advancing this sector's AI adoption rates must be predicated on improvements in these three characteristics for data pipelining and model training/tuning.

Opportunities to Advance:

- Document model training and tuning actions, data sources, and decisions.
- Focus on data storage and sharing options that maintain privacy and security.
- Enhance the ease of user data discovery and retrieval for fine-tuning.
- Focus on advancing methods for incorporating domain-specific knowledge into model responses.
- Promote lower barriers to data publication and sharing where possible.

A.4.2 Data Discovery for Improved Accuracy: Retrieval-Augmented Generation

A popular method for improving LLM response accuracy while decreasing hallucinations, especially for real-time applications, is RAG, which utilizes a pre-trained LLM and then retrieves

additional data to enrich the model's contextual awareness. This method requires a smaller amount of domain-specific data (e.g., organization-provided) than fine-tuning and automatically provides a mechanism to incorporate real-time data into the LLM's output. RAG-enhanced models can decrease model development time and monetary costs.

The RAG pipeline requires that data be indexed and efficiently organized in a vector database[1] so that it can be quickly found and utilized by the LLM after a user or application inputs a prompt. This prompt is then formatted to align better with the vector database. Once the data has been searched and ranked, it can be incorporated into the LLM's response to enhance accuracy and relevance (see Figure A.12).

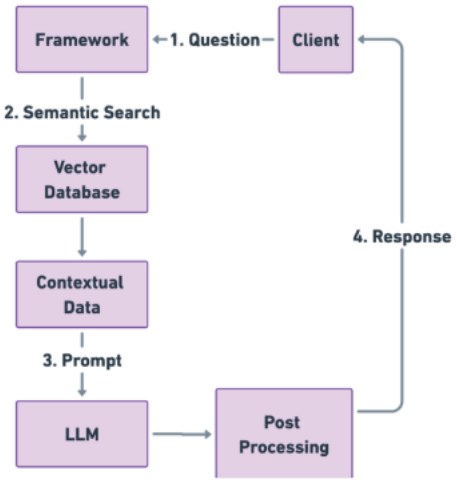

Figure A.12. RAG approach to LLM response generation.[2]

Retrieval-Augmented Fine-Tuning essentially combines the best of both RAG and fine-tuning such that the LLM understands the user's domain thoroughly (i.e., its needs, nuances, and jargon) while also having access to the latest information so it can provide up-to-date, highly relevant answers.[3]

For Emergency Managers

What to Know:

EM organizations should consider using RAG models for domain-specific tasks that require up-to-date or real-time information, such as infrastructure damage assessment and resource allocation during incident response. In the high-consequence environment in which EM operates, these enhanced models will decrease the risks of producing inaccurate or hallucinated LLM outputs while increasing the chances that models provide suggestions based

[1] A vector database stores complex, varied data (audio, video, text, etc.) as numerical representations for quicker retrieval and easier understanding of data relationships.

[2] Parthasarathy, V., Zafar, A., Khan, A., & Shahid, A. (2024). (tech.). *The Ultimate Guide to Fine-Tuning LLMs from Basics to Breakthroughs: An Exhaustive Review of Technologies, Research, Best Practices, Applied Research Challenges and Opportunities* (pp. 6–13). Dublin, Ireland.

[3] Zhang, X., Xie, G., Huang, Y., Xiong, Z., Liu, J., Cui, S., Sun, S., & Sherman Shen, X. (2025). Edge intelligence in the Generative Artificial Intelligence Era. *IEEE Wireless Communications*, 1–9. https://doi.org/10.1109/mwc.2025.3599652

on relevant, recent information. Models that incorporate real-time intelligence will broaden the set of EM tasks, particularly during the response phase, with which an LLM can assist EM practitioners.

Opportunities to Advance:

- Determine the scope of the LLM they would like to deploy (what questions and tasks should it be able to answer) as the primary step.
- Conduct work domain analyses to identify existing tasks relying on real-time information and workflow challenges that might be supported by LLMs and other generative AI models.
- Audit internal databases and other sources of data (notes, emails, etc.) to identify relevant data for the in-scope queries following the determination of scope, including collaborating organization and agency data, which may additionally enhance the LLM's response quality by enabling it to learn patterns related to specific entities, geographies, processes, etc.

For Standards and Policymakers

What to Know:

Due to the rapidly changing nature of incidents, many EM use cases for RAG-based LLMs will require (near) real-time data transmission from edge devices (e.g., sensors, drones, and smartphones) to the AI models themselves for ingestion. Standards and policies will need to be written in a way that facilitates this fast transfer of data without additional barriers that prevent EM users from obtaining up-to-date situational awareness.

Opportunities to Advance:

- Establish data-sharing agreements that clearly outline permissions, scope, timelines, and privacy.
- Establish workshop groups and conduct workshops to fully discover relevant data sources and their time-dependent characteristics (e.g., update frequency).
- Ensure standards support privacy and security while still enabling fast and frequent data transmissions.

For Emergency Management Vendors

What to Know:

Incorporating RAG into LLMs can introduce security threats into the AI systems, including, but not limited to, data poisoning and leakage. Adversaries' actions may vary based on their access model parameters, ability to interact with the model via queries, and knowledge of the

data. Trust is vital to EM tasking, so enhancing the explainability of results to quickly identify the warning signs of adversary behavior will be vital. Further, streaming data comes with quality deviances that will need to be continually addressed as part of the RAG pipeline.

Opportunities to Advance:

- Explainability should be a requirement for models supporting EM to quickly and easily identify when a user should reduce their trust in the model.
- Data should be encrypted during transmission and in storage to prevent tampering.
- Incorporate data updates from edge devices and sensors in addition to using these updates to backfill previously missing or biased information where possible.

For AI Leaders

What to Know:

Real-time data, while necessary for many EM tasks, comes with a set of quality risks, including, but not limited to, missing entries, biases, and errors. Further, the nature of RAG-based models introduces additional data security threat vectors into LLMs. Instituting guardrails can limit potential harm and unintended consequences.

Opportunities to Advance:

- Ensure that data collection for RAG can be continuous and updated as new information becomes available.
- Build data quality assessment and alerting procedures into RAG pipelines.
- Focus research on methods to improve quality in streaming data and LLM abilities to overcome potentially skewed data.
- Build automatic break points into LLM response generation to ensure that models tell users when they do not have enough quality data to generate a trustworthy response (as opposed to hallucinating a response).

A.4.3 Data Quality and Safety Controls for AI Systems

EM AI systems process information from countless sources during incidents- online discourse, citizen photographs, sensor data, third-party feeds, and interagency communications. Not all of this data is reliable, accurate, or appropriate for AI processing. Some may contain degraded information, manipulated content, PII that must be protected, or malicious inputs designed to compromise AI systems. Data quality and safety controls establish protective mechanisms that filter, validate, and sanitize data before it reaches AI models, ensuring that systems only process trustworthy information and that harmful or inappropriate content never influences life-safety decisions. These safeguards operate at multiple points in the data pipeline (see Figure A.13). Content moderation filters screen out harmful media (e.g., graphic violence, inappropriate imagery) before it contaminates training datasets or appears in situational awareness displays. Quarantine pipelines isolate suspect data or processing jobs that exhibit anomalous

characteristics, allowing human review before potentially compromised information affects operations. On-device preprocessing and redaction capabilities strip PII from citizen reports and sensor data at the source, ensuring privacy protection without requiring data to traverse network boundaries where it might be intercepted or misused.

Figure A.13. Data Quality Life Cycle.[1]

For emergency managers, these controls provide confidence that AI recommendations are based on verified, high-quality information rather than contaminated or manipulated data. During major incidents when adversaries may deliberately inject false information or when well-meaning citizens share unverified reports, quality controls prevent AI systems from amplifying incorrect data or making recommendations based on unreliable inputs. These safeguards also protect organizations from legal liability by ensuring that sensitive personal information is properly handled and that AI systems cannot be weaponized through adversarial data injection.

For Emergency Managers

What to Know:

Data quality and safety controls protect AI systems from making recommendations based on faulty data, manipulated content, or compromised data sources during critical incidents. These automated filters and validation mechanisms work continuously to identify suspect information,

[1] StarfishETL. "Data Quality Management: What it is and How to Do it." December 1, 2020. https://www.starfishetl.com/blog/data-quality-management-what-it-and-how-do-it

quarantine potentially harmful inputs, and strip PII before data enters AI processing pipelines. Content moderation ensures that graphic or inappropriate imagery does not contaminate situational awareness systems or training datasets. On-device preprocessing capabilities allow for the collection of valuable information from the public while protecting privacy by removing personal details at the source before data leaves mobile devices or sensors. These controls are especially critical during major incidents when adversaries may deliberately inject false information or when the volume of incoming data makes manual verification impossible.

Opportunities to Advance:

- Implement content moderation filters on all feeds and channels integrated with AI systems, establishing clear policies for what content types are blocked and requiring human review of filtered items.
- Deploy quarantine pipelines that automatically isolate data exhibiting anomalous characteristics (unusual sources, conflicting information, or statistical outliers) for human verification before allowing it to influence AI recommendations.
- Establish on-device redaction capabilities for citizen reporting applications and field sensors, automatically removing PII like names, addresses, and license plates before data is transmitted to operations centers.
- Create escalation protocols requiring human validation before AI systems can act on information from unverified sources, particularly for high-stakes decisions like evacuation orders or resource deployments.
- Conduct regular audits of filtered and quarantined data to identify patterns suggesting coordinated adversarial campaigns or attacks targeting AI systems.

For Standards and Policymakers

What to Know:

Policy frameworks must mandate data quality and safety controls for government AI systems to prevent incorrect data from influencing life-safety decisions and to ensure compliance with privacy regulations. Standards are needed for content moderation in emergency contexts that balance rapid information processing with protection against harmful content. Regulatory requirements should address how PII is handled in AI pipelines, mandating technical controls like on-device redaction rather than relying solely on policy. Liability frameworks must clarify responsibility when AI systems process compromised data that evades quality controls, protecting emergency managers who implement reasonable safeguards while holding accountable those who deploy unprotected systems. Privacy regulations should encourage on-device preprocessing approaches that enable valuable data collection while protecting citizen information at the source.

Opportunities to Advance:

- Establish minimum standards for content moderation and data validation that government AI systems must implement, with certification requirements demonstrating compliance before procurement approval.
- Develop regulatory frameworks mandating on-device redaction of PII for citizen reporting applications and public-facing AI systems, with specific technical requirements and audit procedures.
- Create liability protections for EM agencies that implement documented data quality controls and follow established validation protocols, even when sophisticated attacks evade those safeguards.
- Fund pilot programs testing advanced data quality controls in multi-agency environments, documenting best practices for balancing rapid information processing with protection against harmful content.
- Establish interagency information-sharing protocols that include data quality attestations, allowing receiving agencies to understand what validation and filtering have been applied to shared information.

For Emergency Management Vendors

What to Know:

Building robust data quality and safety controls into AI products differentiates solutions in a market where government customers face increasing threats from adversarial data injection. EM organizations cannot afford AI systems that amplify false information during critical incidents, and they require automated safeguards that operate at the speed and scale of modern data flows. Implementing content moderation, quarantine pipelines, and on-device redaction demonstrates commitment to responsible AI deployment while protecting reputation from incidents where systems process compromised data. These capabilities also address the privacy and security concerns that often delay AI adoption in government, providing technical controls that enable valuable data collection while protecting sensitive information. Vendors who lead in data quality and safety controls position themselves favorably as regulatory requirements mature and as EM organizations learn from incidents where unprotected AI systems propagated faulty data.

Opportunities to Advance:

- Integrate content moderation capabilities into all AI products that process online discourse, citizen reports, or unstructured data, with configurable policies allowing EM customers to define their own filtering criteria.
- Develop quarantine pipeline features that automatically isolate anomalous data for human review, with transparent logging showing what was filtered and why to support accountability and continuous improvement.
- Implement on-device preprocessing and redaction in mobile applications and edge sensors, processing data locally to remove PII before transmission to central systems.

- Create dashboards providing real-time visibility into data quality metrics (sources validated, content filtered, data quarantined), enabling emergency managers to assess information reliability during incidents.
- Establish partnerships with content verification services and threat intelligence providers to enhance a system's ability to detect adversarial attacks targeting EM organizations.

For AI Leaders

What to Know:

EM represents a high-stakes adversarial environment where AI systems must be hardened against data poisoning attacks and privacy violations that could undermine public trust or compromise operations. Data quality and safety controls are not optional features but fundamental requirements for responsible AI deployment in domains where lives depend on information accuracy. Implementing robust content moderation, quarantine pipelines, and on-device preprocessing requires deep technical expertise in adversarial machine learning, privacy-preserving computation, and real-time data validation at scale. These controls must operate with minimal latency—emergency managers cannot wait for lengthy validation processes during rapidly evolving incidents—while maintaining high accuracy to avoid filtering legitimate information that could inform critical decisions. Building these capabilities from the ground up is essential. Retrofitting safety controls into existing architectures often introduces unacceptable performance penalties or security vulnerabilities.

Opportunities to Advance:

- Implement multi-stage content moderation pipelines using both rule-based filters and machine learning classifiers trained on EM contexts, with human-in-the-loop review for edge cases.
- Develop anomaly detection systems that identify suspect data based on source reputation, statistical outliers, temporal patterns, and cross-validation against trusted sources, automatically quarantining questionable information.
- Design on-device preprocessing capabilities using privacy-preserving techniques like federated learning and differential privacy, enabling valuable model training on distributed data without collecting raw PII.
- Create real-time data quality scoring systems that assign confidence levels to incoming information based on source verification, cross-validation, and historical accuracy, surfacing these scores alongside AI recommendations.
- Establish red team testing programs that attempt to degrade data quality and inject adversarial data into AI systems, continuously improving detection capabilities based on attack patterns observed in EM contexts.

A.4.4 Privacy-Preserving Collaboration and Data Sharing

EM inherently requires collaboration across organizational boundaries—multiple agencies, jurisdictions, private sector partners, and even international entities must share information and

coordinate responses during major incidents. However, privacy regulations, security policies, and competitive concerns often prevent organizations from directly sharing the raw data that would enable effective AI-powered coordination. Privacy-preserving collaboration technologies solve this dilemma by enabling organizations to jointly analyze data, train AI models, and generate insights without any party exposing their sensitive raw data to others. These approaches allow EM organizations to leverage collective intelligence while maintaining data sovereignty and compliance with privacy regulations.

Data clean rooms provide secure environments where multiple organizations can analyze combined datasets under strict access controls that prevent any party from extracting raw records belonging to others (see Figure A.14). Organizations contribute encrypted or aggregated data to the clean room, run approved analytical queries or AI models against the joint dataset, and receive only summary results or insights—never the underlying data from partner organizations. This approach enables critical capabilities like regional trend analysis, cross-jurisdictional resource optimization, and multi-agency threat detection while ensuring each organization retains control over its sensitive information and complies with data sharing restrictions.

Figure A.14. Data Clean Room Overview.[1]

On-device preprocessing and redaction capabilities extend privacy protection to the point of data collection, processing information locally on mobile devices, sensors, or edge systems before it ever leaves the source. When citizens report incidents through mobile applications or when sensors collect environmental data, on-device processing can extract valuable features, detect relevant patterns, and strip PII locally. Only the processed, anonymized results are transmitted to central systems, dramatically reducing privacy risks and enabling data collection in scenarios where transmitting raw information would violate regulations or erode public trust. For emergency managers, these privacy-preserving approaches unlock collaboration opportunities and data sources that would otherwise remain unavailable, enabling AI systems to learn from broader datasets while respecting privacy and security boundaries.

[1] Adjust. “How does a data clean room work?” 2025. https://www.adjust.com/glossary/data-clean-room/

For Emergency Managers

What to Know:

Privacy-preserving collaboration technologies enable EM organizations to jointly analyze data and coordinate AI-powered responses with mutual aid partners, other agencies, and private sector entities without violating data sharing restrictions or exposing sensitive information. Data clean rooms allow for combining incident data with regional partners to identify broader patterns, optimize resource allocation across jurisdictions, and train AI models on larger datasets—all while ensuring that no organization can access another's raw records. On-device preprocessing capabilities enable collecting valuable information from citizens and field sensors while protecting privacy by processing data locally and transmitting only anonymized insights to the EOC. These approaches are especially valuable for scenarios like regional disease surveillance, critical infrastructure protection involving private utilities, and cross-border incidents requiring international coordination where direct data sharing would be legally or politically impossible.

Opportunities to Advance:

- Establish data clean room partnerships with regional mutual aid partners to enable joint analysis of incident patterns, resource utilization, and response effectiveness without exposing individual agency records.
- Deploy on-device redaction capabilities in citizen reporting applications and field sensors, processing data locally to extract valuable features while protecting PII before transmission.
- Develop data sharing agreements with private sector partners (e.g., utilities, hospitals, transportation providers) that leverage clean room technologies to enable critical infrastructure coordination while respecting proprietary and regulated data.
- Implement privacy-preserving analytics for after-action reviews involving multiple agencies, allowing joint examination of multi-jurisdictional incident responses without compromising operational security or sensitive information.
- Train staff on the capabilities and limitations of privacy-preserving technologies, ensuring they understand when these approaches enable previously impossible collaboration and when direct data sharing remains necessary.

For Standards and Policymakers

What to Know:

Policy frameworks must evolve to recognize privacy-preserving collaboration technologies as viable alternatives to direct data sharing, updating regulations that currently create barriers to multi-agency AI coordination. Standards are needed to define what constitutes acceptable privacy protection in data clean rooms, on-device processing, and other collaborative approaches to ensure these technologies deliver on their privacy promises. Regulatory guidance should clarify how organizations can use privacy-preserving methods to comply with

data sharing restrictions while still enabling the cross-organizational coordination that EM requires. Liability frameworks must address responsibility when AI systems trained on jointly analyzed data produce errors, allocating accountability among contributing organizations. Investment in privacy-preserving infrastructure (e.g., regional clean rooms, standardized on-device processing frameworks) can dramatically accelerate multi-agency AI adoption by removing the data sharing barriers that currently fragment EM capabilities.

Opportunities to Advance:

- Establish certification standards for privacy-preserving collaboration technologies, defining technical requirements for data clean rooms, on-device processing, and other approaches to ensure they provide genuine privacy protection.
- Develop regulatory guidance explicitly recognizing data clean rooms and privacy-preserving computation as compliant methods for satisfying multi-agency coordination requirements without violating data sharing restrictions.
- Fund regional data clean room infrastructure that EM organizations can use for multi-jurisdictional collaboration, reducing the technical and financial barriers to privacy-preserving coordination.
- Create model data sharing agreements incorporating privacy-preserving technologies, providing templates that EM organizations can adapt for mutual aid partnerships and private sector collaboration.
- Develop liability frameworks for AI systems trained using privacy-preserving multi-party approaches, clarifying how responsibility is allocated when multiple organizations contribute data to joint models.

For Emergency Management Vendors

What to Know:

Privacy-preserving collaboration capabilities represent significant market opportunities as EM organizations seek to unlock cross-agency AI coordination while navigating complex data sharing restrictions. Vendors who provide turnkey data clean room solutions, on-device processing frameworks, and privacy-preserving analytics tools enable customers to participate in regional AI initiatives and public-private partnerships that would otherwise be impossible. These capabilities differentiate products by solving the collaboration problem that often blocks AI adoption in government contexts where privacy regulations, security policies, and interagency politics prevent direct data sharing. Implementing privacy-preserving technologies also builds trust with privacy-conscious EM organizations and improves positioning as regulations increasingly mandate privacy-by-design approaches. Vendors who establish regional clean room platforms or standardized privacy-preserving frameworks can create network effects where solutions become more valuable as adoption grows across agencies.

Opportunities to Advance:

- Develop data clean room capabilities integrated with EM platforms, enabling customers to securely collaborate with partner agencies and private sector entities without exposing raw data.
- Implement on-device processing frameworks in mobile applications and edge sensors, providing privacy-by-design solutions that collect valuable data while protecting PII at the source.
- Create partnership programs connecting EM customers with complementary organizations (e.g., utilities, hospitals, transportation agencies) to enable privacy-preserving cross-sector collaboration using these types of platforms.
- Establish regional clean room infrastructure that multiple EM customers can share, reducing deployment costs and creating network effects as agencies join collaborative analytics environments.
- Provide transparent documentation and audit capabilities demonstrating how privacy-preserving technologies protect data, building trust with security-conscious government customers and supporting their compliance requirements.

For AI Leaders

What to Know:

Privacy-preserving collaboration technologies, including data clean rooms, federated learning, secure multi-party computation, and on-device processing, represent essential capabilities for AI systems operating in the regulated, multi-organizational EM environment. EM organizations possess valuable data for AI training and coordination but face legal, regulatory, and political barriers preventing direct sharing with partners or cloud-based AI providers. Implementing robust privacy-preserving approaches requires expertise in cryptography, distributed systems, and privacy-preserving ML at scale. These technologies must deliver genuine privacy protection while maintaining acceptable performance for real-time emergency operations. On-device processing capabilities demand efficient model architectures and optimization techniques that can operate within the resource constraints of mobile devices and edge sensors. Successfully deploying privacy-preserving AI in EM requires not only technical implementation but also building trust through transparent documentation, third-party audits, and demonstrated compliance with privacy regulations.

Opportunities to Advance:

- Implement data clean room capabilities in AI platforms that enable multiple organizations to jointly train models and generate insights using cryptographic techniques, ensuring no party can access others' raw data.
- Develop on-device AI processing capabilities using techniques like model quantization, pruning, and knowledge distillation to enable privacy-preserving feature extraction and inference on resource-constrained edge devices.

- Create federated learning frameworks allowing EM organizations to collaboratively train AI models by sharing only model updates rather than raw data, with differential privacy guarantees protecting individual records.
- Design transparent privacy accounting systems that document exactly what information is revealed through privacy-preserving collaboration, providing EM customers with evidence they can use to demonstrate regulatory compliance.
- Establish third-party audit programs validating the privacy properties of these systems, building trust with government customers through independent verification rather than relying solely on the technology developer's own security claims.

A.4.5 Real-Time Intelligence and Continuous Model Enhancement

Emergency incidents evolve rapidly—conditions change, fires spread, floods rise, and resource availability fluctuates minute by minute. AI systems trained on historical data alone quickly become outdated during dynamic situations, providing recommendations based on prior conditions rather than the current reality. Real-time intelligence and continuous model enhancement capabilities enable AI systems to automatically incorporate fresh information as it arrives, continuously updating their understanding of evolving situations and refining their predictions based on the latest observations. These streaming pipelines process incoming data from sensors, online discourse, field reports, and other sources in real-time, extracting relevant features and feeding them to AI models so that recommendations reflect current conditions rather than stale snapshots.

Streaming feature pipelines continuously compute and update the inputs that AI models need for decision-making (e.g., outdoor conditions, road closures, facility status, resource locations, and countless other factors that change throughout an incident). Rather than batch-processing data on periodic schedules, streaming pipelines react to events as they occur, ensuring AI systems always have access to the most current information available. Online inference caches and state stores maintain low-latency access to these continuously updated features, enabling AI systems to generate recommendations in seconds rather than waiting for database queries or batch computations. This real-time capability is essential for time-critical EM tasks like routing evacuation traffic, deploying first responders, and predicting hazard spread.

Feature stores provide the architecture that makes real-time intelligence practical, maintaining both online systems optimized for fast access during incidents and offline systems for model training and analysis (see Figure A.15). Time-series feature engineering capabilities extract meaningful patterns from sensor data and event streams (e.g., detecting accelerating trends, identifying anomalies, and computing rolling statistics that capture temporal dynamics). For emergency managers, these real-time capabilities transform AI from a planning tool into an active decision support system that keeps pace with rapidly evolving situations, providing recommendations that reflect the current state of an incident rather than outdated snapshots.

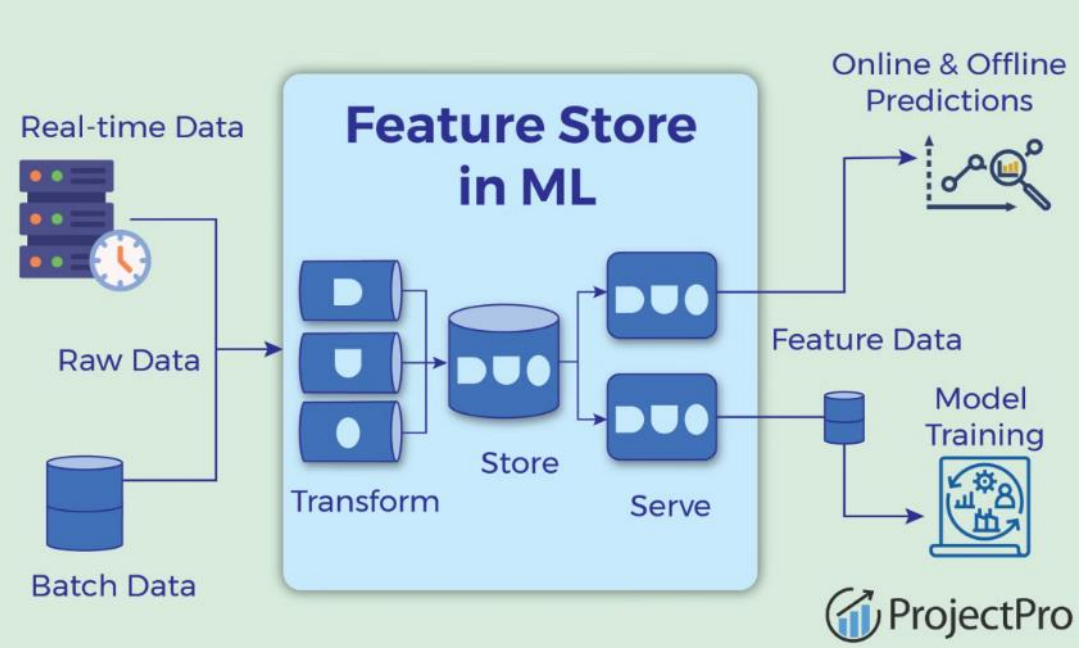

Figure A.15. Feature Store in Machine Learning.[1]

For Emergency Managers

What to Know:

Real-time intelligence and continuous model enhancement capabilities enable AI systems to provide recommendations based on current incident conditions rather than outdated historical data, which is critical for time-sensitive decisions during rapidly evolving emergencies. Streaming feature pipelines continuously process incoming data from sensors, drones, outdoor sensors, and field reports, updating AI models as conditions change throughout an incident. Feature stores maintain both fast-access online systems for live operations and offline systems for training and analysis, ensuring AI recommendations reflect the latest road closures, outdoor conditions, resource positions, and facility status. Time-series feature engineering extracts patterns from sensor data streams (e.g., detecting accelerating fire spread, rising water levels, or degrading infrastructure conditions), enabling predictive capabilities that anticipate incident evolution rather than simply reacting to it.

Opportunities to Advance:

- Implement streaming feature pipelines that continuously process data from edge devices, sensors, and field reports, updating AI systems with current incident conditions rather than relying on periodic batch updates.

[1] ProjectPro. "A Beginner's Guide to Feature Store in Machine Learning." October 28, 2024. https://www.projectpro.io/article/feature-store-in-machine-learning/882

- Deploy feature store architecture, maintaining online caches for fast access during operations and offline stores for model training, ensuring AI recommendations reflect both real-time data and historical patterns.
- Establish time-series monitoring for critical indicators (e.g., fire spread rates, flood levels, traffic flow, resource depletion) with automated alerts when AI detects accelerating trends or anomalous patterns.
- Create experimentation frameworks that safely test real-time AI enhancements in sandboxes before deploying them to production systems, preventing untested changes from affecting live operations during critical incidents.
- Train EOC staff to interpret AI confidence levels and time-sensitivity indicators, ensuring they understand when recommendations reflect fresh data versus when system lags may affect accuracy.

For Standards and Policymakers

What to Know:

Real-time AI capabilities require policy frameworks addressing data quality standards for streaming information, liability for decisions based on continuously updated models, and technical requirements ensuring systems fail safely when data feeds are interrupted. Standards are needed to define acceptable latency for different EM tasks. For example, evacuation routing may require sub-second updates while resource planning can tolerate longer delays. Regulatory guidance must clarify how organizations document decisions when AI models are continuously learning and updating, addressing audit and accountability requirements for systems that do not have static versions in traditional senses. Investment in real-time data infrastructure (e.g., edge processing capabilities, low-latency networks, standardized sensor protocols) can dramatically improve EM AI effectiveness by ensuring fresh data reaches decision-support systems quickly.

Opportunities to Advance:

- Establish data quality standards for streaming information used in real-time AI systems, defining acceptable error rates, latency thresholds, and validation requirements for different EM tasks.
- Develop regulatory frameworks addressing liability when AI systems make recommendations based on continuously updated models, clarifying accountability when rapid data changes affect decision quality.
- Fund edge computing and low-latency network infrastructure enabling real-time data transmission from field sensors, drones, and mobile devices to EOCs and AI systems.
- Create certification programs for real-time AI systems verifying they fail safely when data feeds are interrupted, maintaining human control during system degradation rather than producing unreliable recommendations.
- Establish audit standards for continuously learning AI systems, defining how organizations document model states, data inputs, and recommendation rationales when systems update in real-time rather than through versioned releases.

For Emergency Management Vendors

What to Know:

Real-time intelligence capabilities represent significant market differentiation as EM organizations recognize that AI systems relying solely on historical data cannot keep pace with dynamic incident conditions. Implementing streaming feature pipelines, online inference caches, and continuous model updates requires sophisticated technical architecture but delivers measurable operational value during time-critical scenarios. Feature store capabilities that maintain both fast-access online systems and comprehensive offline stores for training position products as enterprise-grade solutions rather than prototype tools. Time-series feature engineering and anomaly detection capabilities enable predictive AI that anticipates incident evolution, providing emergency managers with early warnings rather than reactive analysis. Vendors who deliver reliable real-time AI with graceful degradation when data feeds fail to build trust in high-stakes operational environments.

Opportunities to Advance:

- Implement streaming feature pipelines in AI products that continuously process incoming sensor data, field reports, and external feeds, updating model inputs as incident conditions change in real time.
- Deploy feature store architecture to maintain online caches optimized for sub-second access during operations and offline stores supporting model training and historical analysis.
- Develop time-series feature engineering capabilities that detect accelerating trends, identify anomalies, and compute rolling statistics from sensor streams, enabling predictive rather than purely reactive AI.
- Create experimentation frameworks allowing EM customers to safely test real-time AI enhancements in sandboxes before production deployment, with feature gates controlling rollout to live operations.
- Build monitoring dashboards showing data freshness, feature update latency, and model confidence levels, giving EOC staff visibility into whether AI recommendations reflect current or stale information.

For AI Leaders

What to Know:

EM represents a demanding real-time AI environment where model recommendations must reflect rapidly changing conditions, including outdoor conditions, infrastructure status, resource availability, and incident evolution. Streaming data comes with quality challenges, including missing entries, sensor errors, network interruptions, and temporal biases that require robust pipeline design and graceful degradation. Feature stores must balance the competing demands of ultra-low latency online serving for operational decisions and comprehensive offline capabilities for model training and experimentation. Time-series feature engineering on

streaming data requires careful design to extract meaningful patterns while avoiding overfitting to transient noise and ensuring features remain stable enough for reliable model performance. Real-time AI systems must fail safely when data feeds are interrupted or degraded, maintaining human control and providing clear indicators of reduced confidence rather than generating unreliable recommendations during system stress.

Opportunities to Advance:

- Design streaming feature pipelines with built-in data quality assessment, automatically detecting missing data, sensor errors, and anomalies while computing confidence scores that reflect input reliability.
- Implement feature store architecture with online systems optimized for ultra-low latency and offline systems supporting efficient batch computation, maintaining consistency between online and offline feature values.
- Develop time-series feature engineering capabilities that extract meaningful patterns from sensor streams while providing robustness to missing data, temporal gaps, and transient noise common in emergency operations.
- Build automatic breakpoints into real-time AI systems that alert users when data quality or freshness degrades below thresholds, ensuring models communicate uncertainty rather than generating confident but unreliable recommendations.
- Create comprehensive monitoring and observability for streaming pipelines, tracking data freshness, processing latency, feature drift, and model performance to enable rapid diagnosis when real-time capabilities degrade during critical incidents.

A.5 Memory and Personalization

Memory and personalization maintains short- and long-term contextual knowledge, user preferences, and experiential learning that allow agents to adapt over time. By separating storage and retrieval mechanisms from specific databases or models, it supports future evolutions in memory technology.

A.5.1 Handling Missing and/or Degraded Data Quality

AI models, including LLMs, usually function best when using data of a quality similar to the data on which they were trained. When pieces of data are missing, or the data quality is degraded, such as due to increased uncertainty, the accuracy of the AI model can decrease rapidly.

The impacts of a disaster will likely degrade the quality of existing data streams and increase the proportion of missing data. For example, physical damage to water pipelines and infrastructure could disrupt pressure, flow, and leak-detection sensors, causing erratic or incomplete data. Seismic impacts might displace sensor installations, leading to misaligned or false readings. Satellite signals used for remote monitoring might encounter delays or obstruction from debris or atmospheric disturbances. Additionally, water contamination caused by ruptured pipelines could introduce anomalies in water quality data collected by sensors.

In a situation with severely degraded data quality, a data enhancement engine (a platform that cleans and improves data quality before feeding data streams to the decision support LLM) could leverage statistical models that include encoding dependencies between data streams during normal conditions. When data quality is degraded, these models could be used to detect

anomalies or inverted to impute small amounts of missing or degraded data. This could allow the decision support model to continue functioning as well as possible in the presence of uncertain data without requiring additional data sources.

For Emergency Managers

What to Know:

In disaster and EM scenarios, degraded data quality or missing data will almost certainly lead to decreased performance from AI decision support models. Systems should be integrated into the decision support pipeline to handle this problem.

Opportunities to Advance:

- Identify multiple useful data streams measuring a similar quantity, for redundancy.
- Assess (qualitatively) acceptable uncertainty levels, both in the final output of a decision support model, and in the raw measurements being collected.

For Standards and Policymakers

What to Know:

Policymakers may need to set standards for the maximum uncertainty that decision support models and data enhancement engines can gracefully handle to ensure that vendors and AI leaders develop trustworthy products.

Opportunities to Advance:

- Evaluate the highest level of uncertainty under which EM operations can still function.

For Emergency Management Vendors

What to Know:

While a variety of statistical methods exist for handling missing or degraded data, and for assessing data dependencies and redundancies, no off-the-shelf products provide these services in a format appropriate for a non-technical end user. AI companies may be traditionally less concerned with these issues.

Opportunities to Advance:

- Create a software platform or dashboard that evaluates and mitigates missing data and degraded data issues, appropriate for use by a non-technical user.

For AI Leaders

What to Know:

In an EM scenario, AI models will need to gracefully handle multiple modes: a normal operations mode and a disaster recovery mode. Data availability and quality will shift rapidly. This could include high rates of missing data or a sudden increase in uncertainty, causing the data distribution to be drastically different, but confidence in the decision support model must remain the same. In these scenarios, redundant data streams can be an asset.

Opportunities to Advance:

- Develop emergency-specific techniques for handling missing data and high uncertainty; these should account for the possibility of leveraging redundant or dependent data streams.
- Investigate techniques for handling missing data and sudden shifts in uncertainty within the LLM itself.

A.5.2 Handling Missing Data Modalities

AI models and LLMs usually rely on a specific, pre-defined set of data modalities to make their predictions (e.g., water flow rate, traffic cameras). Usually, this is the same as the set of modalities on which the model was originally trained. If one or more of these modalities suddenly becomes available, the quality of the LLM's predications will degrade or the LLM may be completely unable to function beyond raising an alert that one or more required modalities are missing.

For example, camera imagery may become completely unavailable during an earthquake due to infrastructure damage. In this case, if the decision support model signals that this modality is missing but crucial for the model to function, the data enhancement engine could use the results of previously completed analyses to check for a sufficient redundant or proxy data stream to send to the decision support model in place of the missing modality. This could include previously available data streams detected via a correlation analysis, identifying strong linear or non-linear relationships between the two streams over time, such as consistent patterns in water pressure and flow rate data from the same pipelines. The engine could also leverage causal inference techniques, analyzing whether one data stream reliably predicts the behavior of another, marking it as a suitable proxy. Additionally, the model might use feature importance analysis to determine that one stream contributes less unique information than another. By continuously ingesting and analyzing real-time data, the engine could learn temporal changes or environmental contexts where redundancy arises, storing that information for future use during a disaster.

Dealing with a missing modality could also include identifying a previously unavailable data stream with a similar modality to replace the missing data. This could be a data stream that becomes available due to the disaster recovery process, which involves creative re-purposing of data collection resources available in the area. For example, if imagery has become unavailable, a rural area might have farm drones available that could be re-purposed to collect useful imagery related to disaster recovery. For a discussion about quickly investigating new data sources, see Section 5.6.

For Emergency Managers

What to Know:

If a decision support model suddenly loses access to one or more data modalities, rather than those modalities being available but degraded, the model likely will not function properly. It will be able to raise an alert that crucial data is unavailable, but it may not be able to proceed without regaining access.

Opportunities to Advance:

- Proactively identify proxy or redundant data modalities in their locations. Creative re-purposing of other data sources should be considered, such as leveraging drones for imagery if normally available cameras are offline.

For Standards and Policymakers

What to Know:

Policymakers should either proactively ensure that data sharing across organizations is possible with minimal overhead or be prepared to navigate data sharing across organizations as smoothly as possible in real time.

Opportunities to Advance:

- Develop policies and procedures for data sharing across organizations, specifically for emergency scenarios. Consider that some of this data might be coming from private citizens.

For Emergency Management Vendors

What to Know:

While technical methods exist for identifying proxy data streams, platforms for performing this operation and replacing data modalities in a way that is autonomous or accessible to non-technical users do not exist. The agility required in an EM scenario is unique, and the speed at which a potentially large amount of data may need to be moved likely exceeds normal operations for a general-purpose AI model.

Opportunities to Advance:

- Develop a software platform that can identify proxy data modalities and gracefully replace a missing modality, making sure to include flexibility to ingest entirely new data streams.

For AI Leaders

What to Know:

In an EM scenario, entire data modalities may become unavailable, and the AI decision support model must handle this gracefully. Ideally, the accuracy of the model would remain the same even when lacking access to some data streams.

Opportunities to Advance:

- Ensure that AI models can raise an alert when crucial data is missing and performance is no longer trustworthy. Consider having the model also describe what crucial data is required.
- Ensure that an AI model's robustness to missing modalities is fully characterized.

A.5.3 Control Protocols for Sudden Retroactive Access to Data

AI models and LLMs, especially those that leverage any temporal data, usually assume that data is moving forward in time, sometimes at a pre-set frequency, or within a pre-set time window size. Models are not usually designed to incorporate sudden access to past data and would not usually be able to leverage this type of data access.

For example, damage to infrastructure such as power or network connections could cause a partial or complete data blackout for a period of time. When connections are restored, assuming some or all the data collected during the blackout was stored locally on devices, the data enhancement engine and the decision support model would suddenly have retroactive access to all monitoring data. For example, if a disaster took out data connections for 24 hours, as soon

as service was restored the data enhancement engine would have access to all data from the previous 24 hours in addition to new real-time data. This is likely a change from how the data enhancement engine and decision support model are originally designed to function; however, the information in the retroactive data could be crucial. It may provide additional detail about what occurred during the disaster or more information about the extent of the damage.

The data enhancement engine could have a mode to handle a sudden influx of past data in several ways:

- Data Gap Analysis and Backfilling: The data enhancement engine could identify gaps in the data streams caused by the outage and compare them to real-time operational data collected since connectivity was restored. By aligning timestamps, it can fill the missing intervals using the retroactive data, ensuring a continuous dataset for analysis.
- Anomaly Detection: The data enhancement engine could detect anomalies or trends that may have occurred during the outage (e.g., sudden spikes in pipeline leaks or rapid drops in reservoir levels). This retrospective analysis could be used to adjust the decision support model to account for missed events.
- Recalibration and Historical Contextualization: With access to retroactive data, the engine could recalibrate itself, such as its tracking of redundant data streams. Historical patterns and correlations could also be revalidated, allowing for improved quality of data fed to the decision support model.
- Priority-Based Processing: The engine would need to prioritize incorporating critical data streams (such as those related to infrastructure damage or water contamination) that directly impact disaster recovery efforts. This ensures the decision support model can provide the most critical insights derived from the retroactive data first.
- Proxy or Synthetic Data Alignment: If the enhancement engine used proxy or synthetic data during the outage to estimate missing parameters, it could now compare those estimates with the actual retroactive data, ensuring correction of inaccuracies and improving confidence levels in its previous calculations.

For Emergency Managers

What to Know:

Incorporating data from the past may be difficult for AI models, despite the utility of that data.

Opportunities to Advance:

- Be prepared to identify when retroactive data may become available; for example, when network connections will be restored to sensors.

For Standards and Policymakers

What to Know:

Access to past data that becomes available when infrastructure is restored after a disaster could provide crucial information to decision support. Standards and policies should be in place to ensure that this data is collected and accessible.

Opportunities to Advance:

- Develop policies and standards for storing as much sensed data as possible locally and at the edge to maximize the information stored during a network outage.

For Emergency Management Vendors

What to Know:

Techniques are required for maintaining data collection during infrastructure outages and for later incorporating this data into an AI model.

Opportunities to Advance:

- Maximize the ability to sense and store data locally and at the edge in the case of network or power outages. This may require a significant amount of storage availability within the EM infrastructure.
- Create a data enhancement engine platform that gracefully handles sudden access to past data, quickly analyzing the newly available data for anomalies and backfilling the existing data stream.

For AI Leaders

What to Know:

In an EM scenario, the AI model may suddenly gain access to a large amount of past data. This could be due to a network connection being restored or many other situations. AI models must be prepared to incorporate data from the past at unpredictable points in their pipeline.

Opportunities to Advance:

- Develop AI methods for gracefully incorporating past data into current inferences quickly; the suddenly available past data may be a very large dataset.

A.5.4 Multi-Jurisdictional Knowledge Access and Compliance

EM operations frequently span multiple jurisdictions (i.e., cities, counties, states, and federal agencies), each with its own data sovereignty requirements, privacy regulations, and operational boundaries. When an AI system needs to answer questions or make recommendations during a multi-jurisdictional incident like a regional wildfire or hurricane, it must access relevant knowledge from multiple organizational sources while respecting legal and policy boundaries that restrict where data can be stored, who can access it, and how it can be used. Traditional centralized knowledge bases that collect all information in one location violate these requirements, creating legal risks and preventing many organizations from participating in AI-enhanced coordination.

Federated RAG and jurisdiction-partitioned indices solve this challenge by allowing AI systems to query knowledge across organizational boundaries without centralizing sensitive data (see Figure A.16). Rather than copying all EM plans, facility information, and resource inventories into a single database, federated approaches maintain data in separate indices controlled by each jurisdiction. When an AI system needs information to answer a question, such as available hospital capacity across a region or mutual aid resources that can respond to an incident, it queries multiple local indices simultaneously, retrieves relevant documents from each jurisdiction, and synthesizes a response. Hierarchical RAG architectures add layers of specificity, allowing AI to first search personal notes, then team knowledge, then organizational databases, and finally public information, retrieving context appropriate to the user's role and clearance level.

Jurisdiction-partitioned indices go further by ensuring that retrieval operations comply with data residency and access control requirements. For example, when an AI system serves users in California, it queries indices physically located in California and containing only data that California regulations permit to be accessed. When the same system serves federal emergency managers, it can access additional indices containing information that state-level users cannot see. This architecture allows AI systems to operate across complex regulatory environments, supporting both routine operations where access rules are well-defined and disaster scenarios where emergency data sharing agreements temporarily expand access, while maintaining audit trails proving that every retrieval operation complied with applicable policies.

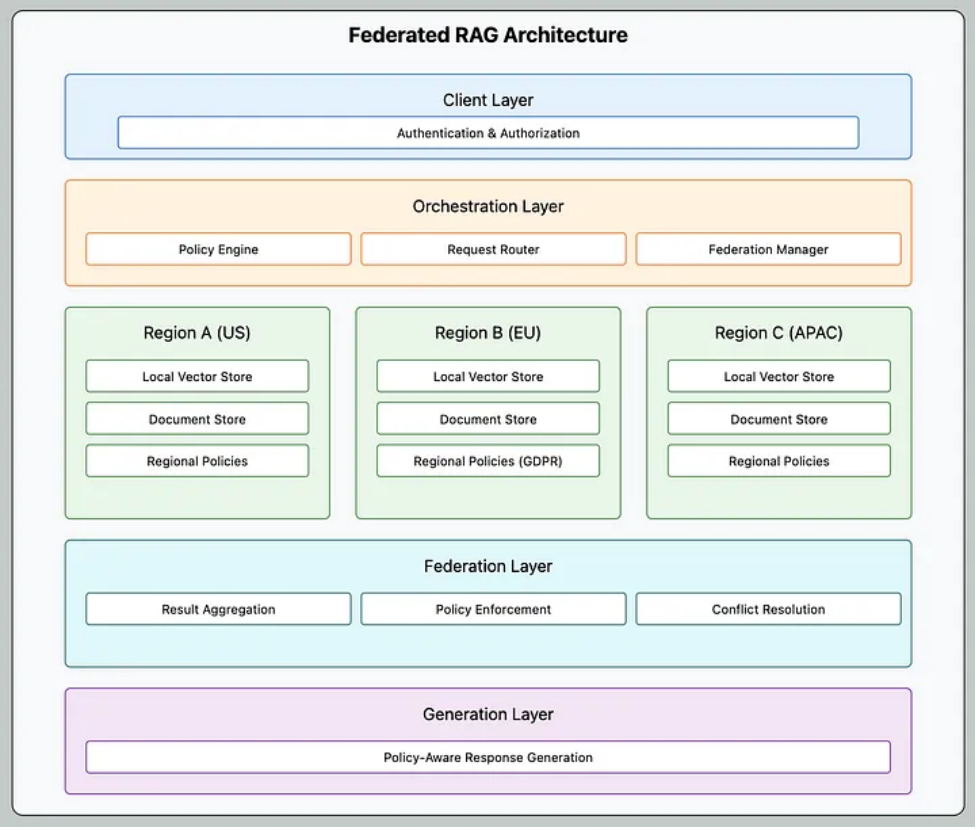

Figure A.16. Federated RAG Architecture.[1]

For Emergency Managers

What to Know:

Multi-jurisdictional knowledge access architectures enable AI systems to retrieve relevant information from multiple agencies and jurisdictions during regional incidents while respecting data sovereignty, privacy regulations, and operational security requirements that prevent centralizing sensitive information. Federated RAG allows AI tools to query mutual aid partners' resource inventories, neighboring jurisdictions' facility databases, and state/federal planning documents without requiring those organizations to share raw data or consolidate information in external systems. Hierarchical retrieval ensures that AI responses incorporate personal notes, team knowledge, and organizational context appropriate to each user's role and clearance level rather than providing the same generic answers to everyone. Jurisdiction-partitioned indices maintain compliance by ensuring AI retrieval operations respect data residency requirements and access control policies, automatically limiting what information each user can access based on their location, organization, and authorities.

[1] Nigam, G. "A Complete Guide to Implementing Federated RAG." December 16, 2024. https://medium.com/aingineer/a-complete-guide-to-implementing-federated-rag-671bbda94e5a

Opportunities to Advance:

- Implement federated RAG capabilities that allow AI systems to query partner organizations' knowledge bases during multi-jurisdictional incidents, establishing data sharing agreements that define what information can be retrieved and under what circumstances.
- Deploy hierarchical retrieval architectures that layer personal, team, organizational, and public knowledge sources, ensuring AI recommendations incorporate context appropriate to each user's role while protecting sensitive operational information.
- Establish jurisdiction-partitioned indices for a given organization's knowledge that can be selectively exposed to mutual aid partners based on data sharing agreements, maintaining fine-grained control over what external users can retrieve.
- Configure AI systems to maintain audit trails of all cross-jurisdictional retrieval operations, documenting what information was accessed from which organizations to support after-action reviews and compliance verification.
- Train EOC staff on how federated and hierarchical retrieval affects AI responses, ensuring they understand that recommendations may differ based on their access level and that partners may receive different information based on data sharing boundaries.

For Standards and Policymakers

What to Know:

Policy frameworks must evolve to support federated knowledge architectures that enable AI-enhanced multi-jurisdictional coordination while maintaining data sovereignty and compliance with privacy regulations. Standards are needed to define how jurisdictions expose knowledge for federated retrieval, what metadata must accompany retrieved documents to prove compliance, and how audit trails document cross-organizational access. Regulatory guidance should clarify how emergency data sharing agreements interact with federated AI systems, addressing whether temporary access expansions during disasters apply to automated retrieval operations or require human approval. Liability frameworks must address responsibility when AI systems synthesize information from multiple jurisdictions and produce recommendations that span organizational boundaries. Investment in federated knowledge infrastructure (e.g., regional indices, standardized retrieval protocols, compliance verification systems) can dramatically improve multi-agency AI coordination while respecting data sovereignty requirements that currently fragment EM capabilities.

Opportunities to Advance:

- Develop regulatory frameworks explicitly recognizing federated knowledge architectures as compliant methods for multi-jurisdictional AI coordination, clarifying how they satisfy data sharing requirements while maintaining sovereignty.
- Establish standards for jurisdiction-partitioned indices defining technical requirements for access control, audit logging, and compliance verification that prove retrieval operations respect data residency and privacy regulations.

- Create model data sharing agreements for federated RAG that define what types of knowledge can be retrieved across organizational boundaries, under what circumstances, and with what audit requirements.
- Fund regional federated knowledge infrastructure that EM organizations can use for multi-jurisdictional AI coordination, reducing technical and financial barriers to compliant cross-organizational retrieval.
- Develop liability frameworks for AI systems that synthesize information from multiple federated sources, clarifying how responsibility is allocated when recommendations span jurisdictional boundaries.

For Emergency Management Vendors

What to Know:

Federated and hierarchical knowledge access capabilities represent significant market opportunities as EM organizations seek AI coordination across jurisdictional boundaries while navigating data sovereignty requirements. Vendors who provide federated RAG solutions, jurisdiction-partitioned indices, and hierarchical retrieval architectures enable customers to participate in regional AI initiatives without violating data residency regulations or exposing sensitive information to external systems. These capabilities differentiate products by solving the multi-jurisdictional coordination problem that often blocks AI adoption in government contexts where centralized knowledge bases violate sovereignty requirements. Implementing compliant federated architectures also builds trust with privacy-conscious and security-focused EM organizations. Vendors who establish regional federated knowledge platforms or standardized retrieval protocols can create network effects where solutions become more valuable as adoption grows across jurisdictions.

Opportunities to Advance:

- Develop federated RAG capabilities that allow AI systems to query multiple organizational knowledge bases simultaneously, synthesizing responses while maintaining audit trails proving each retrieval operation's compliance with data sharing agreements.
- Implement hierarchical retrieval architectures that automatically layer personal, team, organizational, and public knowledge sources based on user roles and clearance levels, ensuring appropriate context without manual configuration.
- Create jurisdiction-partitioned index infrastructure that maintains data residency compliance by physically locating indices in appropriate regions and enforcing access controls based on user location and authorities.
- Establish regional federated knowledge platforms that multiple EM customers can join, reducing deployment costs and creating network effects as agencies participate in shared retrieval networks.
- Provide transparent audit and compliance reporting showing what information was retrieved from which jurisdictions for each AI response, enabling customers to verify that operations respected data sharing boundaries.

For AI Leaders

What to Know:

EM represents a complex federated environment where AI systems must retrieve knowledge across organizational and jurisdictional boundaries while respecting data sovereignty, access control, and compliance requirements that prevent centralization. Implementing robust federated RAG, hierarchical retrieval, and jurisdiction-partitioned architectures requires expertise in distributed systems, query federation, access control, and compliance verification at scale. These systems must operate with acceptable latency—emergency managers cannot wait for lengthy cross-organizational queries during time-critical incidents—while maintaining strong security and audit capabilities. Hierarchical retrieval architectures must intelligently balance relevance across knowledge layers, avoiding situations where personal notes overwhelm organizational knowledge or where public information dilutes specialized expertise. Jurisdiction-partitioned indices demand careful design, ensuring that compliance rules are enforced at the retrieval layer rather than relying on application-level checks that can be bypassed. Successfully deploying federated knowledge systems in EM requires not just technical implementation but also building trust through transparent audit trails and demonstrated compliance with data sovereignty requirements.

Opportunities to Advance:

- Implement federated RAG architectures that query multiple organizational indices in parallel, synthesizing responses while maintaining ultra-low latency acceptable for operational decision support during emergencies.
- Develop hierarchical retrieval systems that intelligently weight and blend results from personal, team, organizational, and public knowledge layers based on relevance, recency, and user context rather than simple concatenation.
- Design jurisdiction-partitioned indices with compliance enforcement at the storage and retrieval layers using cryptographic access controls, physical data residency, and immutable audit logs proving every operation's regulatory compliance.
- Create transparent explainability interfaces showing which jurisdictions and knowledge layers contributed to each AI response, enabling emergency managers to assess information provenance and verify cross-organizational retrieval compliance.
- Establish performance monitoring for federated retrieval operations tracking query latency, index availability, and cross-jurisdictional access patterns to ensure systems maintain acceptable responsiveness during multi-agency incidents.

A.5.5 Contextual Knowledge Integration for Enhanced Decision Support

AI systems for EM must integrate multiple types of knowledge (such as structured databases, unstructured documents, spatial relationships, temporal patterns, and expert domain knowledge) to provide accurate and contextually appropriate recommendations. A purely keyword-based search might retrieve documents mentioning "shelter" but miss the critical context that certain facilities are unsuitable after an earthquake, that others have limited capacity, or that some are already at capacity based on real-time data. Similarly, a purely semantic search using vector embeddings might retrieve conceptually similar documents but fail

to capture the precise legal requirements, geographic constraints, or organizational relationships that determine what actions are permissible and feasible during an incident.

Hybrid knowledge architectures combining vector databases, knowledge graphs, and multiple retrieval strategies solve these challenges by enabling AI systems to understand both the semantic meaning and the structured relationships within EM domains. Vector databases enable fast similarity search over embeddings, quickly finding documents conceptually related to a query, even when they do not share exact keywords. Knowledge graphs capture structured domain knowledge: the hierarchical relationships between jurisdictions, the dependencies between critical infrastructure systems, the organizational command structures that determine decision authority, and the temporal sequences that define proper response procedures. Hybrid approaches fuse these capabilities, using knowledge graphs to constrain and guide vector retrieval, ensuring AI systems find semantically relevant information that also satisfies the structural requirements of EM operations.

For emergency managers, these hybrid knowledge systems enable AI to provide recommendations that are not only topically relevant but also operationally feasible, legally compliant, and contextually appropriate. When an AI system recommends shelter locations, it retrieves documents about available facilities (vector search), filters them based on geographic proximity and disaster type (knowledge graph), checks current capacity against real-time data, and verifies that recommended actions comply with interagency agreements and regulatory requirements. This integration of multiple knowledge types transforms AI from a document retrieval tool into a decision support system that understands the complex operational context in which EM decisions must be made.

For Emergency Managers

What to Know:

Contextual knowledge integration architectures enable AI systems to provide recommendations that account for the complex relationships, constraints, and requirements that govern EM operations, going beyond simple document retrieval to understand organizational hierarchies, geographic dependencies, temporal sequences, and regulatory requirements. Hybrid systems combining vector databases and knowledge graphs allow AI to quickly find conceptually relevant information while ensuring recommendations respect operational realities like facility capacities, jurisdictional boundaries, mutual aid agreements, and command authorities. These architectures are critical for multi-faceted EM tasks where decisions must simultaneously satisfy multiple constraints. Hybrid knowledge integration ensures AI recommendations are not just topically relevant but operationally feasible and contextually appropriate for a specific situation.

Opportunities to Advance:

- Implement hybrid knowledge systems that combine vector databases for fast semantic search with knowledge graphs encoding critical relationships like organizational hierarchies, facility dependencies, resource constraints, and jurisdictional boundaries.

- Develop or procure knowledge graphs capturing EM domain structures, including ICS roles, NIMS resource typing, mutual aid agreements, critical infrastructure dependencies, and regulatory requirements that constrain operations.
- Configure AI systems to use knowledge graphs as constraints on vector retrieval, ensuring recommendations satisfy structural requirements like geographic proximity, organizational authority, and resource availability rather than just conceptual relevance.
- Establish processes for continuously updating knowledge graphs with operational changes (e.g., new facilities, modified mutual aid agreements, updated resource inventories, changed organizational structures), ensuring AI recommendations reflect current realities.
- Train EOC staff on how hybrid knowledge integration affects AI recommendations, helping them understand that responses incorporate both semantic relevance and operational constraints encoded in structured knowledge.

For Standards and Policymakers

What to Know:

Policy frameworks must recognize that effective AI for EM requires integration of multiple knowledge types, not just document repositories, necessitating standards for knowledge graphs encoding organizational relationships, regulatory requirements, and operational constraints. Standards are needed to define how EM domain knowledge should be structured (e.g., NIMS resource typing, ICS organizational hierarchies, facility categorizations, hazard classifications) to enable knowledge graph interoperability across jurisdictions. Regulatory guidance should address how organizations maintain and update knowledge graphs as operational realities change, ensuring AI systems do not make recommendations based on outdated organizational structures or obsolete mutual aid agreements. Investment in standardized EM knowledge graphs capturing common domain structures that all jurisdictions share can dramatically accelerate AI adoption by providing ready-made frameworks rather than requiring each organization to build domain knowledge representations from scratch.

Opportunities to Advance:

- Establish standards for EM knowledge graphs defining how to encode common domain structures, including NIMS resource types, ICS organizational roles, facility classifications, hazard categories, and interagency relationships.
- Develop certification programs for EM AI systems, verifying they incorporate structured knowledge graphs constraining retrieval and recommendations, not just semantic document search capabilities.
- Fund development of standardized EM ontologies that jurisdictions can adopt and extend, providing common frameworks for representing domain knowledge rather than requiring each organization to build from scratch.
- Create regulatory requirements for AI systems to maintain provenance showing how recommendations incorporate both retrieved documents and structured knowledge constraints, enabling audit of decision rationale.

- Establish update and maintenance standards for operational knowledge graphs, defining how organizations document changes to mutual aid agreements, facility status, organizational structures, and other dynamic information affecting AI recommendations.

For Emergency Management Vendors

What to Know:

Hybrid knowledge integration capabilities combining vector databases and knowledge graphs represent significant market differentiation as EM organizations recognize that purely semantic document retrieval cannot capture the complex operational constraints governing their decisions. Vendors who provide integrated solutions (such as pre-built EM ontologies, tools for maintaining knowledge graphs, and hybrid retrieval engines that combine semantic and structural search) enable customers to deploy AI that understands operational context rather than just document similarity. Implementing domain-specific knowledge graphs demonstrates deep EM expertise and builds trust with customers who have experienced AI systems making technically correct but operationally infeasible recommendations. These capabilities also create opportunities for value-added services, helping organizations develop their knowledge graphs, maintaining ontologies as standards evolve, and integrating knowledge from multiple jurisdictions for regional coordination.

Opportunities to Advance:

- Develop hybrid knowledge systems combining vector databases for semantic document search with knowledge graphs encoding EM domain structures, including organizational hierarchies, facility relationships, resource constraints, and regulatory requirements.
- Provide pre-built EM ontologies aligned with NIMS, ICS, and other standards, allowing customers to deploy knowledge graphs without building domain representations from scratch while supporting customization for local needs.
- Create user-friendly tools for EM organizations to maintain and update knowledge graphs as operational realities change, ensuring AI recommendations reflect current facility status, mutual aid agreements, organizational structures, and resource availability.
- Implement explainability interfaces showing how hybrid retrieval combined semantic document search with knowledge graph constraints to produce recommendations, helping emergency managers understand why certain options were included or excluded.
- Establish partnerships with EM standards organizations and domain experts to continuously improve ontologies and knowledge representations, ensuring systems incorporate evolving best practices and regulatory requirements.

For AI Leaders

What to Know:

EM represents a domain where effective AI requires integration of heterogeneous knowledge types (e.g., unstructured documents, structured databases, semantic embeddings, symbolic relationships, temporal patterns, and spatial constraints) that must be queried and reasoned over simultaneously. Implementing robust hybrid knowledge architectures combining vector databases and knowledge graphs requires expertise in information retrieval, graph databases, query federation, and constraint satisfaction at scale. These systems must intelligently balance semantic similarity from vector search with structural constraints from knowledge graphs, avoiding situations where conceptually relevant but operationally infeasible recommendations dominate results. Knowledge graph maintenance presents unique challenges in EM, where operational realities change frequently. Mutual aid agreements are updated, facilities are damaged or repaired, and organizational structures evolve. These types of changes require systems that gracefully handle incomplete or temporarily inconsistent knowledge. Successfully deploying hybrid knowledge systems requires not just technical implementation but also deep domain expertise to encode meaningful EM relationships and constraints that improve decision quality.

Opportunities to Advance:

- Implement hybrid retrieval architectures that combine vector similarity search with knowledge graph traversal, using symbolic relationships to constrain, filter, and re-rank semantically relevant documents based on operational feasibility and compliance requirements.
- Develop EM knowledge graphs encoding critical domain structures, including organizational hierarchies, facility dependencies, resource constraints, geographic relationships, temporal sequences, and regulatory requirements that govern operations.
- Create query planning systems that intelligently balance semantic and structural search based on query characteristics, using knowledge graphs heavily for structured questions and vector search for open-ended queries.
- Design knowledge graph maintenance workflows supporting incremental updates without requiring complete regeneration, enabling EM organizations to continuously reflect operational changes like facility status updates and organizational modifications.
- Establish monitoring and quality assurance for hybrid retrieval operations, tracking how often knowledge graph constraints improve recommendation quality versus introducing false negatives, continuously tuning the balance between semantic relevance and structural correctness.

A.6 Governance and Operations

Operations and governance provides deployment oversight, monitoring, compliance, security, and policy enforcement across the entire stack. It provides guardrails, auditing, and life-cycle management to ensure responsible, reliable, and cost-effective operation independent of specific Development Operations (DevOps) or governance tooling.

A.6.1 Real-Time Data Ingestion Services

AI models are usually designed to make inferences based on data modalities they are already expecting. It may be difficult to incorporate access to a new data stream if the new data is not in a format expected by the model. AI models for EM must be designed to incorporate unexpected data streams and to ingest these data streams quickly and seamlessly.

Part of the disaster preparation activities should include that formatting requirements for all usually monitored data modalities are clearly documented, and this documentation should be stored locally in case of connectivity outages. Having this documentation easily accessible will assist with ingesting additional data sources that become available during disaster recovery. For example, after an earthquake disrupts water supplies in a town and first responders start arriving, a new data modality that could become available is real-time geolocated reports from mobile devices or communication systems used by first responders. These reports could include GPS-tagged data such as photos, videos, or text detailing burst pipelines, damaged water tanks, or contaminated water sources. Drones used by response teams could also provide aerial imagery of the impacted zones, identifying blocked access points or areas where water distribution has been cut off. Additionally, with the influx of emergency response vehicles, traffic and logistical data from local transportation networks might become important. For these new data sources to be useful, their outputs must be organized and formatted in a way that the larger decision support model can use. Using the documented formatting requirements as part of an LLM prompt, an LLM could be used to generate the required adapters to seamlessly ingest these new data streams for use in the decision support model.

In addition, during a disaster and recovery, online discourse from the disaster area may be crucial for situational awareness and decision support. In some cases, it may be key in detecting the highest priority response areas when other data modalities are degraded or unavailable. Data from online discourse can pose a challenge for analysis due to its unstructured nature, potential propriety aspects depending on the specific online communication platform, and concerns around privacy. Online discourse data would need to be filtered quickly and reduced to only signals relevant to disaster recovery. This could be accomplished via prompting an LLM to generate queries for the online communication platform API specifically relevant to the disaster. For example, filtering the online discourse data stream down to only instances of civilians with geotags near the area of interest mentioning water. The LLM could then also be used to generate code to preprocess this filtered online discourse media data stream into a cleaned time series useful to the decision support model.

For Emergency Managers

What to Know:

EM organizations should be aware of the potential to incorporate new data streams into decision support models during recovery operations. As additional personnel arrive on scene, they may also be equipped to carry sensors that contribute to situational awareness.

Opportunities to Advance:

- Equip first responders with digital sensors, recording devices, etc. that can communicate with EM infrastructure.

For Standards and Policymakers

What to Know:

Incorporating new data streams on the fly during a disaster will require data formatting standards and policies for data sharing across organizations.

Opportunities to Advance:

- Develop standards for data formatting in an EM scenario, in collaboration with technical experts.
- Develop policies for sharing data across organizations, potentially including industry (e.g., online discourse) and private citizens.

For Emergency Management Vendors

What to Know:

EM models must be able to quickly ingest new data modalities that become available during disaster recovery. This may require building data adapters in real time.

Opportunities to Advance:

- Develop agile technology, likely using an LLM, to create data adapters that properly format data in real time.
- Develop technology, appropriate for non-technical users, to scrape online discourse for information relevant to a specific EM scenario.
- Work with policymakers to understand data formatting requirements and standards.

For AI Leaders

What to Know:

During a public safety event, entirely new data streams or modalities may become available due to additional personnel or technologies arriving on scene. The formatting of these datasets may be arbitrary, and AI decision support models must handle this gracefully.

Opportunities to Advance:

- Work with policymakers to develop acceptable standards for data formatting for AI models.
- Develop AI techniques that are robust to differing data formats or small errors in data formatting.

A.6.2 Data Sovereignty and Cross-Jurisdictional Compliance

EM operations frequently involve sensitive data that must comply with complex geographic and jurisdictional restrictions. Privacy laws vary by state and country, certain data cannot leave specific regions, and multi-agency coordination must respect each organization's sovereignty over its information. When AI systems process emergency data, they must ensure that information stays within legally permitted boundaries, that processing occurs in authorized locations, and that data movements comply with applicable regulations. Data residency routing and enforcement capabilities automatically direct data and AI processing to compliant geographic locations, preventing violations that could expose organizations to legal liability or compromise interagency trust.

Egress control and data loss prevention (DLP) systems monitor and restrict how data leaves AI systems and organizational boundaries, ensuring that sensitive information (i.e., personally identifiable details about citizens, proprietary infrastructure data, or classified operational intelligence) cannot be inadvertently shared, exfiltrated by adversaries, or leaked through AI-generated outputs (see Figure A.17). These controls become especially critical when AI systems integrate data from multiple jurisdictions or when emergency operations involve public-private partnerships where different parties have varying rights to access and use shared information. Control and data plane separation enables centralized policy management while enforcing restrictions locally, allowing EM organizations to coordinate regional AI initiatives while maintaining sovereignty over their data.

For emergency managers, these capabilities enable multi-jurisdictional AI coordination without violating data sovereignty requirements or creating compliance risks. During regional incidents spanning multiple states or international borders, AI systems can analyze combined intelligence while ensuring each jurisdiction's data remains within permitted boundaries. These architectural controls provide the technical enforcement that transforms data sharing agreements and mutual aid memorandums from aspirational policies into operational reality, enabling the cross-organizational collaboration that effective EM requires while respecting the legal and political boundaries that govern data handling.

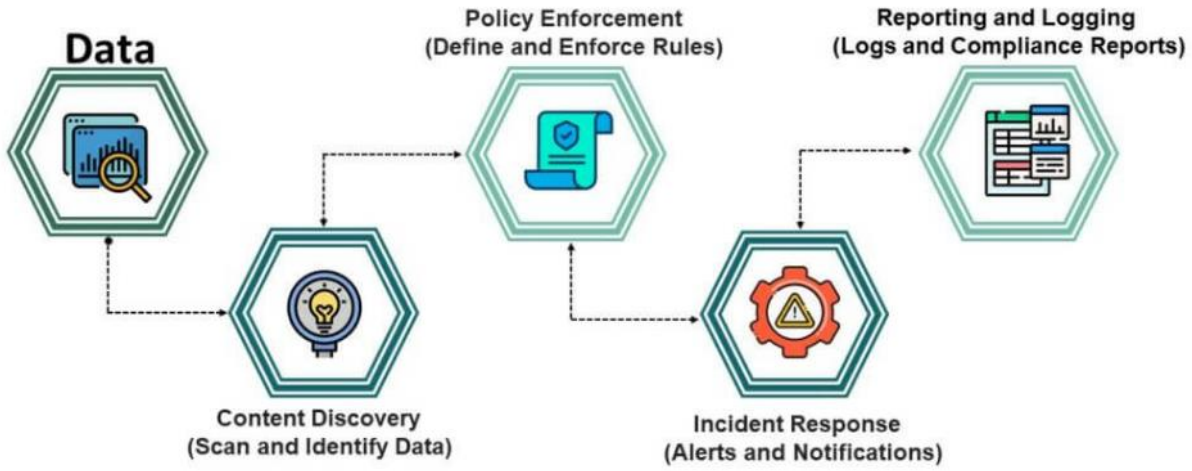

Figure A.17. Data Loss Prevention.[1]

For Emergency Managers

What to Know:

Data sovereignty and compliance controls enable EM organizations to participate in regional AI initiatives and multi-jurisdictional emergency coordination while maintaining control over where sensitive data is stored, processed, and shared. Data residency routing ensures that when AI systems process citizen information, critical infrastructure data, or operational intelligence, that processing occurs within legally permitted geographic boundaries, preventing violations of state privacy laws, federal data protection requirements, or interagency agreements. Egress controls and DLP prevent sensitive information from inadvertently leaving systems through AI-generated reports, recommendations shared with partner agencies, or data breaches. Separation of control and data planes allows a given organization to participate in regional coordination platforms and shared AI services while maintaining local enforcement of specific data handling policies, ensuring centralized collaboration does not compromise sovereignty over sensitive information.

Opportunities to Advance:

- Implement data residency routing in AI systems processing sensitive information, configuring geographic boundaries that ensure data storage and processing comply with applicable state, federal, and interagency restrictions.
- Deploy egress controls and DLP monitoring all AI outputs (e.g., reports, recommendations, shared intelligence) to prevent inadvertent disclosure of PII, classified data, or information restricted by data sharing agreements.

[1] Kumar, G. “Data Loss Prevention (DLP).” 2025. https://www.educba.com/data-loss-prevention/

- Establish separation between control planes (where policies are defined and coordination occurs) and data planes (where sensitive data is stored and processed), enabling participation in regional AI platforms without centralizing data.
- Document data sovereignty requirements in procurement specifications for AI systems, requiring vendors to demonstrate compliance with geographic restrictions and providing audit capabilities proving where data was stored and processed.
- Train EOC staff on data sovereignty implications of AI system use, ensuring they understand geographic restrictions affecting what data can be shared with which partner agencies during multi-jurisdictional incidents.

For Standards and Policymakers

What to Know:

Policy frameworks must address how AI systems comply with complex, overlapping jurisdictional requirements governing EM data, including state privacy laws, federal regulations, international agreements, and interagency memorandums that restrict where data can be stored, processed, and shared. Standards are needed that define how data residency is verified and enforced in AI systems, what audit trails prove compliance with geographic restrictions, and how control/data plane separation maintains sovereignty while enabling coordination. Regulatory guidance should clarify how emergency data sharing agreements interact with AI systems that may process information across jurisdictional boundaries, addressing whether automated processing constitutes sharing and what controls satisfy compliance requirements. Liability frameworks must address responsibility when AI systems inadvertently violate data sovereignty restrictions, protecting organizations that implement reasonable controls while holding accountable those who ignore jurisdictional requirements.

Opportunities to Advance:

- Develop regulatory frameworks explicitly defining data residency requirements for EM AI systems, clarifying which data types must remain within specific geographic boundaries and what technical controls satisfy compliance.
- Establish certification standards for AI systems verifying they enforce data sovereignty restrictions, including requirements for audit trails proving where data was stored and processed during emergency operations.
- Create model data sharing agreements for multi-jurisdictional AI coordination that explicitly address control/data plane separation, defining how organizations can participate in shared platforms while maintaining sovereignty over their information.
- Fund development of regional AI infrastructure that provides compliant data residency for EM operations, reducing the technical and financial burden of organizations building their own geographically distributed systems.
- Develop liability protections for EM organizations implementing documented data sovereignty controls, even when sophisticated attacks or system failures cause inadvertent violations.

For Emergency Management Vendors

What to Know:

Data sovereignty and compliance capabilities represent essential requirements—not optional features—for AI systems serving multi-jurisdictional EM markets where complex geographic restrictions govern data handling. Vendors who provide robust data residency routing, egress controls, and control/data plane separation enable customers to participate in regional AI initiatives and mutual aid coordination without violating sovereignty requirements that currently fragment EM capabilities. These capabilities differentiate solutions in government markets where procurement specifications increasingly mandate compliance with state privacy laws, federal data protection requirements, and interagency agreements. Implementing transparent audit capabilities that prove where data was stored and processed builds trust with security-conscious customers and supports their compliance obligations. Vendors who establish regional infrastructure supporting compliant data residency can create network effects where specified platforms become the preferred solution as adoption grows across jurisdictions.

Opportunities to Advance:

- Implement data residency routing in AI platforms, allowing customers to configure geographic boundaries and ensuring data storage and processing comply with applicable jurisdictional restrictions, with automated enforcement preventing policy violations.
- Deploy comprehensive egress controls and DLP monitoring for all AI outputs, preventing inadvertent disclosure of sensitive information while providing transparency into what data was shared, with whom, and under what authority.
- Design control/data plane separation architectures enabling centralized policy management and coordination while maintaining customer sovereignty over their data, allowing participation in regional platforms without centralization.
- Provide detailed audit capabilities documenting where every piece of data was stored, processed, and transmitted, with immutable logs proving compliance with jurisdictional restrictions during emergency operations.
- Establish regional infrastructure in key geographic areas supporting compliant data residency for EM customers, reducing their deployment complexity while creating competitive advantages through network effects.

For AI Leaders

What to Know:

EM represents a complex regulatory environment where AI systems must navigate overlapping jurisdictional requirements governing data storage, processing, and movement across city, county, state, federal, and sometimes international boundaries. Implementing robust data sovereignty and compliance capabilities requires expertise in distributed systems, policy enforcement, geographic routing, and audit logging at scale. These systems must make

real-time decisions about where data can be processed and shared based on complex, context-dependent rules that vary by data type, user role, operational phase, and applicable regulations. Data residency routing must operate transparently without introducing unacceptable latency. Egress controls must balance security with operational necessity, preventing data leakage while allowing legitimate information sharing during multi-jurisdictional emergency response. Successfully deploying compliant AI systems requires not just technical implementation but also deep understanding of EM legal frameworks, interagency agreements, and regulatory requirements that govern data handling.

Opportunities to Advance:

- Implement data residency routing that automatically directs data storage and processing to compliant geographic locations based on data classification, user location, and applicable regulations, with real-time enforcement preventing policy violations.
- Develop egress controls using content analysis, contextual rules, and user permissions to prevent unauthorized data disclosure while allowing legitimate sharing, with comprehensive logging documenting every information transfer and its authorization.
- Design control/data plane architectures that separate policy management from data processing, enabling centralized coordination while maintaining distributed data sovereignty with cryptographic controls ensuring separation is enforced.
- Create transparent audit capabilities generating immutable logs proving where data was stored, processed, and transmitted, with searchable interfaces enabling EM customers to verify compliance with jurisdictional restrictions.
- Establish monitoring systems detecting data sovereignty violations in real-time, with automated alerting and remediation capabilities ensuring rapid response when compliance failures occur during emergency operations.

A.6.3 Accountability, Audit, and Incident Response for AI Systems

EM organizations face legal obligations to document their decisions, maintain records for investigations, and demonstrate compliance with regulations governing data handling and AI use. When AI systems support life-safety decisions such as recommending evacuations, allocating resources, or predicting hazard spread, organizations must be able to prove what data informed those recommendations, how AI models processed that information, and whether proper procedures were followed. Data retention, archival, and deletion policies ensure that organizations maintain appropriate records without creating unnecessary legal exposure from outdated or irrelevant information. Audit logging and e-discovery capabilities provide immutable records of AI system operations, enabling after-action reviews, regulatory investigations, and legal proceedings to reconstruct exactly what happened during critical incidents.

Incident response and kill switch capabilities provide emergency controls when AI systems malfunction, are compromised by adversaries, or begin producing dangerous recommendations. During a cybersecurity incident, organizations need the ability to immediately disable AI access to sensitive data, shut down compromised systems, or revert to manual operations without waiting for lengthy approval processes. Kill switches prevent cascading failures where AI recommendations based on poisoned data or compromised models propagate across interconnected EM systems. These controls become especially critical in multi-agency

environments where one organization's compromised AI could affect partners' operations if not rapidly isolated.

For emergency managers, these accountability and incident response capabilities provide the confidence to deploy AI systems in high-stakes operational environments. Comprehensive audit trails protect organizations during legal proceedings by documenting that decisions followed proper protocols and were based on the best available information. Retention policies ensure critical records are preserved while avoiding the legal liability of maintaining unnecessary data indefinitely. Incident response capabilities provide the safety net that allows aggressive AI adoption—knowing that systems can be rapidly disabled if they malfunction or are compromised reduces the risk of deployment while maintaining operational flexibility.

For Emergency Managers

What to Know:

Accountability and incident response controls provide the documentation and emergency capabilities necessary to safely deploy AI systems in high-stakes EM environments where decisions have legal consequences and system failures could endanger lives. Audit logging maintains immutable records of what data AI systems accessed, what recommendations they produced, and what actions were taken based on those recommendations. Data retention and archival policies ensure critical records are preserved according to legal requirements while systematically deleting outdated information that creates unnecessary legal exposure. Incident response and kill switch capabilities allow immediate disablement of AI systems when they malfunction, are compromised, or produce dangerous recommendations, preventing cascading failures and enabling rapid reversion to manual operations during emergencies.

Opportunities to Advance:

- Implement comprehensive audit logging for all AI system operations, recording what data was accessed, what models produced which recommendations, what confidence levels accompanied outputs, and what actions users took based on AI guidance.
- Establish data retention policies aligned with legal requirements and organizational needs, defining how long different types of AI-related records must be preserved and implementing automated deletion of data that no longer serves operational or legal purposes.
- Deploy incident response capabilities including kill switches that allow immediate disablement of AI systems or specific AI features when they malfunction or are compromised, with clear protocols defining who has authority to activate emergency controls.
- Create e-discovery processes enabling rapid retrieval of AI-related records during investigations, with indexed, searchable audit logs supporting reconstruction of decision-making processes during critical incidents.
- Conduct regular incident response drills that test kill switch procedures and manual operation fallbacks, ensuring EOC staff can rapidly disable AI systems and maintain operations without AI support during emergencies.

For Standards and Policymakers

What to Know:

Policy frameworks must define audit, retention, and incident response requirements for EM AI systems to ensure accountability while avoiding requirements so burdensome they prevent beneficial AI adoption. Standards are needed specifying what AI operations must be logged, how long different types of records must be retained, what audit trail capabilities enable effective investigations, and what incident response controls must be available. Regulatory guidance should clarify how existing records retention requirements apply to AI systems, addressing whether AI recommendations constitute official records and what documentation proves compliance with decision-making protocols. Liability frameworks must address responsibility when organizations properly document AI-assisted decisions versus when they fail to maintain adequate records, protecting those who implement comprehensive accountability measures. Investment in standardized audit frameworks and incident response tools can accelerate AI adoption by providing ready-made compliance capabilities rather than requiring each organization to develop custom solutions.

Opportunities to Advance:

- Establish minimum audit logging requirements for EM AI systems, defining what operations must be recorded, what metadata must accompany logs, and how long audit records must be retained.
- Develop data retention standards aligned with existing records requirements, clarifying how AI-related information should be categorized, what constitutes a "record" requiring preservation, and what can be systematically deleted.
- Create certification programs for AI systems verifying they include incident response capabilities meeting EM needs, including kill switches, degraded operation modes, and manual override procedures.
- Fund development of standardized e-discovery tools enabling EM organizations to efficiently search and retrieve AI-related records during investigations without requiring expensive custom solutions.
- Develop liability protections for organizations implementing documented accountability measures, clarifying that proper audit trails and incident response capabilities provide legal safe harbor even when AI systems produce errors.

For Emergency Management Vendors

What to Know:

Accountability and incident response capabilities represent essential requirements for AI systems serving EM markets where legal obligations, regulatory oversight, and operational safety demand comprehensive documentation and emergency controls. Vendors who provide robust audit logging, automated retention management, and well-designed incident response capabilities enable customers to confidently deploy AI while satisfying legal requirements and

maintaining operational safety. These capabilities differentiate solutions in government markets where procurement specifications increasingly mandate accountability measures and where customers have learned from incidents where inadequate documentation created legal liability. Implementing user-friendly e-discovery interfaces and automated compliance reporting reduces the operational burden on resource-constrained EM organizations. Vendors who establish industry-leading accountability practices build trust and may influence emerging regulatory standards.

Opportunities to Advance:

- Implement comprehensive audit logging capturing all AI system operations with immutable, tamper-evident records, including what data was accessed, what models produced which outputs, user actions, and system configurations at the time of operation.
- Develop automated retention management applying configurable policies defining how long different types of AI-related data are preserved and implementing systematic deletion, ensuring compliance without creating unnecessary legal exposure.
- Design incident response capabilities including granular kill switches that allow selective disablement of specific AI features or complete system shutdown, with clear authorization controls and automated notifications when emergency measures are activated.
- Create user-friendly e-discovery interfaces enabling EM staff to search and retrieve AI-related records without technical expertise, supporting rapid response to investigation requests and legal proceedings.
- Provide compliance dashboards showing audit coverage, retention policy compliance, and incident response readiness, giving EM customers confidence that their accountability measures satisfy regulatory requirements.

For AI Leaders

What to Know:

EM represents a highly regulated, legally consequential environment where AI systems must provide comprehensive accountability through immutable audit trails, policy-driven retention management, and robust incident response capabilities. Implementing these controls requires expertise in secure logging, tamper-evident storage, policy engines, and fail-safe architectures that maintain safety during system degradation. Audit logging must capture sufficient context to reconstruct decision-making processes—not just what recommendations AI produced but what data informed them, what confidence levels accompanied outputs, and what alternatives were considered. Retention management must balance legal preservation requirements with privacy regulations requiring systematic deletion, implementing automated policies that adapt as data ages and regulatory requirements change. Incident response capabilities must fail safely, ensuring kill switches do not create new vulnerabilities and that degraded operation modes maintain essential functionality while isolating compromised components.

Opportunities to Advance:

- Implement immutable audit logging using cryptographic techniques, ensuring records cannot be tampered with after creation and capturing comprehensive context, including data provenance, model versions, confidence scores, and user interactions.
- Develop policy-driven retention management that automatically classifies AI-related data, applies appropriate preservation periods, and implements systematic deletion, ensuring compliance with both preservation requirements and privacy regulations.
- Design multi-level incident response capabilities, including feature-specific kill switches, degraded operation modes maintaining essential functionality, and complete system disablement, with authorization controls preventing unauthorized activation.
- Create explainability and reconstruction capabilities enabling investigators to replay AI decision-making processes using archived data, model snapshots, and audit logs, proving what information was available and how systems processed it.
- Establish monitoring and alerting for accountability system health, detecting when audit logging fails, retention policies are not executing, or incident response capabilities degrade, ensuring EM customers maintain continuous compliance.

A.6.4 Trust and Quality Assurance for AI Operations

AI systems deployed in EM must maintain trustworthy operation even as the data they process changes, the models they use evolve, and the operational environments they support shift over time. Data quality and validation pipelines continuously assess incoming information, detecting errors, anomalies, and suspicious patterns before they contaminate AI recommendations. When AI systems make life-safety decisions based on degraded sensor data, manipulated online discourse, or corrupted databases, the consequences can be catastrophic. Quality controls provide the first line of defense, ensuring AI processes only verified, reliable information.

Model and data drift monitoring detects when the statistical properties of incoming data diverge from what AI models were trained on, or when model performance degrades as operational conditions change. During major incidents, data distributions can shift dramatically (e.g., outdoor conditions become extreme, infrastructure behaves abnormally, populations move in unexpected ways), causing AI models trained on historical norms to produce unreliable recommendations. Drift detection (Figure A.18) alerts emergency managers when model confidence should be reduced, when manual verification should increase, or when systems should revert to human decision-making. Model registry and versioning capabilities maintain comprehensive records of which AI models were deployed when, what data they were trained on, and how they performed, enabling organizations to trace decisions back to specific model versions during after-action reviews and to roll back to previous versions when new models underperform.

CI/CD for data and models enables systematic testing and controlled rollout of AI improvements while maintaining operational safety. Rather than deploying untested AI enhancements directly into production systems during emergencies, organizations can validate changes in controlled environments, gradually roll them out using canary or blue-green deployment strategies, and monitor performance before full adoption. For emergency managers, these trust and quality assurance capabilities provide confidence that AI recommendations are based on reliable data and proven models, that degrading performance will be detected quickly, and that the complete history of AI system evolution is documented for accountability and continuous improvement.

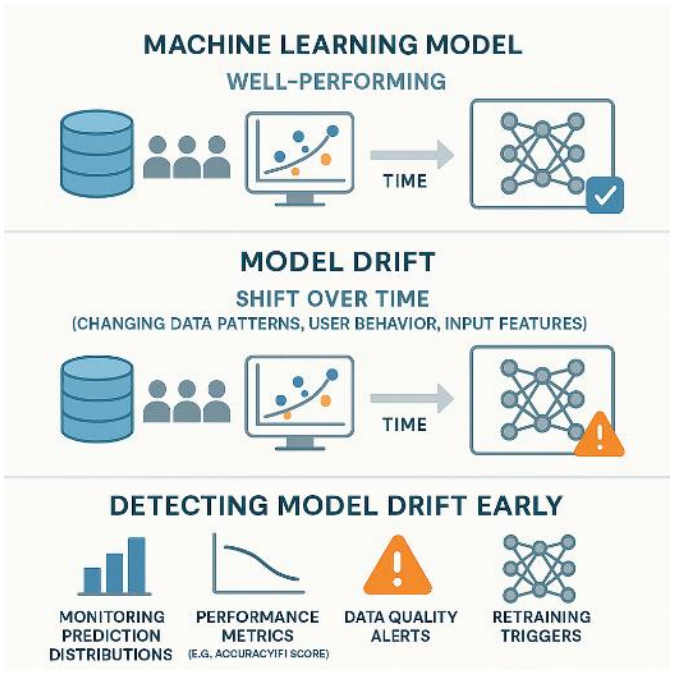

Figure A.18. Model Drift Detection.[1]

For Emergency Managers

What to Know:

Trust and quality assurance controls ensure AI systems maintain reliable performance even as data quality fluctuates during incidents and as models evolve over time. Data quality and validation pipelines continuously assess incoming information, detecting sensor errors, anomalous readings, and suspicious patterns before they influence AI recommendations. Model and data drift monitoring alerts when the statistical properties of incoming data diverge from training conditions or when AI performance degrades, signaling when to reduce reliance on AI and increase human oversight. Model registry and versioning maintain complete records of which AI models were deployed during each incident, enabling after-action reviews to trace decisions to specific model versions and supporting rollback when new models underperform. CI/CD capabilities enable systematic testing of AI improvements in controlled environments before production deployment, preventing untested changes from affecting live operations during critical incidents.

Opportunities to Advance:

- Implement data quality and validation pipelines that continuously assess incoming sensor data, field reports, and external feeds, automatically flagging anomalous readings, missing data, and suspicious patterns before they reach AI models.

[1] Tech Sandesh. "What Is Model Drift and How to Detect It Early." June 3, 2025. https://techsandesh.in/2025/06/03/what-is-model-drift-and-how-to-detect-it-early/

- Deploy model and data drift monitoring that compares current data distributions and model performance against baseline conditions, generating alerts when divergence exceeds acceptable thresholds and degraded AI reliability requires increased human oversight.
- Establish model registry and versioning systems that maintain comprehensive records of which AI models were active during each incident, what data they were trained on, and how they performed, supporting after-action reviews and enabling rollback to previous versions.
- Create CI/CD processes for testing AI improvements in sandbox environments before production deployment, using canary or blue-green strategies to gradually roll out changes while monitoring performance before full adoption.
- Train EOC staff to interpret data quality indicators, drift alerts, and model confidence scores, ensuring they understand when AI recommendations are reliable versus when degraded conditions require reverting to manual decision-making.

For Standards and Policymakers

What to Know:

Policy frameworks must define minimum quality assurance standards for EM AI systems to ensure trustworthy operation while avoiding requirements so prescriptive they prevent beneficial innovation. Standards are needed specifying what data quality metrics must be monitored, what levels of model drift trigger alerts or automatic disablement, what documentation model registries must maintain, and what testing processes must precede production deployment. Regulatory guidance should clarify how organizations demonstrate that AI systems maintained acceptable performance during incidents, addressing what evidence proves compliance with quality assurance requirements. Policymakers must set standards for maximum uncertainty that AI systems can gracefully handle, ensuring vendors develop products that alert users when data quality or model confidence degrades below operational thresholds. Investment in standardized quality assurance frameworks can accelerate AI adoption by providing ready-made compliance capabilities rather than requiring each organization to develop custom solutions.

Opportunities to Advance:

- Establish minimum data quality monitoring requirements for EM AI systems, defining what metrics must be tracked, what validation checks must be performed, and what quality thresholds trigger alerts or automatic restrictions.
- Develop standards for model drift detection specifying acceptable divergence levels from training conditions, when alerts must be generated, and when AI systems must automatically reduce confidence or disable features until human review occurs.
- Create certification programs for AI systems verifying they include comprehensive model registries documenting training data, performance characteristics, deployment history, and version lineage meeting EM accountability requirements.

- Define CI/CD requirements for EM AI, including mandatory testing protocols, staged rollout procedures, and performance monitoring before full production deployment of model updates.
- Evaluate the highest level of data uncertainty and model drift under which EM operations can still function safely, establishing regulatory limits that AI systems must respect.

For Emergency Management Vendors

What to Know:

Trust and quality assurance capabilities represent essential requirements for AI systems serving EM markets where unreliable recommendations during critical incidents could endanger lives and expose organizations to legal liability. Vendors who provide robust data quality monitoring, drift detection, comprehensive model registries, and systematic CI/CD processes enable customers to confidently deploy AI while maintaining operational safety and regulatory compliance. These capabilities differentiate solutions in government markets where procurement specifications increasingly mandate quality assurance measures. Implementing user-friendly dashboards that communicate data quality status, model confidence levels, and drift alerts in operationally meaningful terms reduces the technical burden on EM organizations. Creating platforms that evaluate and mitigate data quality issues appropriate for non-technical users addresses a critical gap where statistical methods exist but accessible products do not.

Opportunities to Advance:

- Implement comprehensive data quality and validation pipelines with user-friendly dashboards showing real-time quality metrics, validation results, and alerts when incoming data exhibits errors, anomalies, or suspicious patterns.
- Develop model and data drift monitoring with configurable thresholds that generate alerts when performance degrades or data distributions diverge, with clear explanations helping emergency managers understand when to reduce AI reliance.
- Create robust model registries maintaining complete documentation of training data sources, performance characteristics, deployment history, and version lineage, with interfaces enabling rapid lookup during incidents and after-action reviews.
- Establish CI/CD platforms supporting systematic testing in sandbox environments, staged rollout using canary or blue-green deployment, and automated performance monitoring before full production release of AI improvements.
- Build software platforms or dashboards that evaluate and mitigate missing data and degraded data issues, appropriate for use by non-technical EM users who need quality assurance without requiring statistical expertise.

For AI Leaders

What to Know:

EM represents a demanding environment where AI systems must maintain trustworthy operation despite rapidly changing data quality, sudden distribution shifts during disasters, and the life-critical consequences of unreliable recommendations. Implementing robust quality assurance requires expertise in statistical process control, drift detection, anomaly identification, and graceful degradation under uncertainty. Data quality monitoring must operate in real-time with minimal latency while distinguishing genuine quality issues from the legitimate but extreme data patterns that occur during major incidents. Drift detection must account for EM's unique characteristic that disasters fundamentally change data distributions. In other words, what appears as drift requiring model shutdown may be the disaster scenario the model should handle. Model registries must comprehensively document not just model artifacts but training data provenance, performance across different scenarios, and known limitations. CI/CD processes must balance rapid innovation with operational safety, enabling systematic testing without delaying critical improvements during evolving threat landscapes.

Opportunities to Advance:

- Implement real-time data quality monitoring using statistical methods that detect missing entries, sensor errors, biases, and anomalies while accounting for the legitimately extreme values that occur during disasters, with automatic alerts when quality degrades below operational thresholds.
- Develop drift detection systems that distinguish between distribution shifts requiring model caution (training data no longer representative) versus shifts the model should handle (disaster conditions it was designed for), with intelligent alerting that considers operational context.
- Create comprehensive model registries documenting training data sources, validation performance across different scenarios, including disaster conditions, known limitations and failure modes, deployment history, and complete version lineage supporting accountability and rollback.
- Establish CI/CD pipelines supporting automated testing in environments replicating emergency conditions, staged rollout with performance monitoring, and rapid rollback capabilities when new models underperform during actual incidents.
- Build automatic breakpoints into AI systems, ensuring models alert users when data quality or confidence degrades below trustworthy thresholds, explicitly stating what crucial data is missing or degraded rather than generating potentially unreliable recommendations.

Appendix B – Lessons Learned and Best Practices from AI Operationalization

This section details the lessons learned and best practices from previous operational AI tool deployments that the EMOTR team has deemed to be the most impactful. To frame this discussion, it is critical to note that public safety and EM outcomes are improved when timely, accurate, relevant, and parsimonious data is available. The collection, aggregation, indexing, storage, sharing, and retrieval of such information forms the unseen majority of effort required to enable the desired data management, sharing, and retrieval (DMSR) architecture. This highly complex mission engages a range of stakeholders across public and private sectors. While several foundational frameworks may apply, such as maximal expected utility[1] and minimum description length,[2] no consensus exists on the best design or principles for creating such a DMSR.

The need for a DMSR is becoming critical, as the scale and velocity of data available to public safety officials and emergency managers are increasing,[3] the availability of data-driven tools provides unprecedented opportunities, but also incredible risks,[4] and the scale and frequency of public safety threats increase. As in many complex systems, an *a priori* holistic design is unrealistic, if not impossible, and would become obsolete shortly after its implementation. However, organizing principles and retrospective analysis of successful and unsuccessful examples of these principles (best practices and lessons learned, respectively) can elucidate useful localized rules, the application of which leads to an improved overall DMSR. This section first discusses two frameworks which, while developed in other fields, can be usefully applied to public safety DMSR. Thereafter, this section discusses three best practices and lessons learned: two straightforward examples and a third that may be either a best practice or lesson learned (depending on the details). These examples may be useful guidance to both public and private entities as nodes within the DMSR are developed and ultimately lead to a better DMSR.

B.1 Frameworks

The concept of maximal expected utility from economics[1] applies directly to the acquisition, sharing, and retrieval of data and to the fields of public safety and EM. Other fields have applied expected utility to information, notably the field of medical research[5] coining the term Value of Information. Defining expected utility in a public safety context requires the enumeration of loss of life, injuries, and economic damage across varying future timescales and under sizable uncertainty. In public safety and EM, utility is primarily defined by cost—both in monetary cost, and in the cost of lives and suffering of the public—as compared to any “return”, which does not have any obvious definition in public safety or EM. The expectation must be applied over both uncertainty and future timescales: information has utility proportional with how surely it can be

[1] R. A. Briggs, “Normative Theories of Rational Choice: Expected Utility,” in *The Stanford Encyclopedia of Philosophy*, E. N. Zalta and U. Nodelman, Eds., Winter 2023.Metaphysics Research Lab, Stanford University, 2023.

[2] J. Rissanen, “Modeling by shortest data description,” *Automatica*, vol. 14, no. 5, pp. 465–471, Sep. 1978, doi: 10.1016/0005-1098(78)90005-5

[3] S. Misra, P. Roberts, and M. Rhodes, “Information overload, stress, and emergency managerial thinking,” *International Journal of Disaster Risk Reduction*, vol. 51, p. 101762, Dec. 2020, doi: 10.1016/j.ijdrr.2020.101762

[4] Y. Zhu *et al.*, “A Survey of Data Agents: Emerging Paradigm or Overstated Hype?” arXiv, Oct. 2025. doi: 10.48550/arXiv.2510.23587. Available: https://arxiv.org/abs/2510.23587. [Accessed: Nov. 19, 2025]

[5] K. Claxton, “The irrelevance of inference: A decision-making approach to the stochastic evaluation of health care technologies,” *Journal of Health Economics*, vol. 18, no. 3, pp. 341–364, Jun. 1999, doi: 10.1016/S0167-6296(98)00039-3, H. W. Tuffaha, L. G. Gordon, and P. A. Scuffham, “Value of information analysis in healthcare: A review of principles and applications,” *Journal of Medical Economics*, vol. 17, no. 6, pp. 377–383, Jun. 2014, doi: 10.3111/13696998.2014.907170

used to minimize future cost to the public. This uncertainty also helps avoid "future-bias," where far-future outcomes have an undue effect on the expected utility. While multifaceted, a holistic analysis of the expected utility or value of information for each subsystem or component in a DMSR can provide principled guidance where otherwise it would be impossible.

Another useful framework is the concept of minimum description length from information theory.[1] Minimum description length states that the minimum length of data required to convey the information is optimal. Applying this to DMSR in public safety is critical to avoid the information overload that is presently lamented in public safety and EM. It is also naturally counterweighted by the expected utility: while the maximal expected utility prefers saving all possible data, minimum description length prefers saving minimal *information.*[2] For example, a list of the status of all water mains in a municipality is data, whereas the location of a malfunctioning water main is information. Minimum description length would prefer a readout of only the single malfunctioning water main to a readout of the status of all.

Using a balanced analysis of the above two frameworks, it becomes possible to perform a principled analysis of DMSR strategies and designs. Applying these frameworks to several DMSR components and subsystems yielded several important or counter-intuitive results below.

B.2 Lesson Learned – Data Unavailability

The most disruptive lesson learned identified was inaccessible information. Public[3] and private[4] entities have decried data silos. Further, a U.S. Executive Order was recently published to abolish data silos.[5] Many dimensions exist that can make data inaccessible: sensitivity can keep it from end users without proper need-to-know, system interconnection can inhibit it from transmission to end users, and undiscoverability can prevent its request. Evidence is emerging that centralization/corporatization can also cause data inaccessibility through anticompetitive request denials.[6] All of these pressures for data inaccessibility are present, and in fact dominant, in public safety and EM. While the ethics involved, especially around data sensitivity[7], are nuanced and not easily dismissed, inaccessible data has zero utility, despite its acquisition and

---

[1] J. Rissanen, "Modeling by shortest data description," *Automatica*, vol. 14, no. 5, pp. 465–471, Sep. 1978, doi: 10.1016/0005-1098(78)90005-5, P. Grünwald and T. Roos, "Minimum Description Length Revisited," *arXiv.org*. https://arxiv.org/abs/1908.08484v2, Aug. 2019. doi: 10.1142/S2661335219300018

[2] Information theory differentiates data from information whereby the latter is only those elements of the former which surprise the received.

[3] R. Rodd, "DOE Data Days 2025 Report," Lawrence Livermore National Laboratory (LLNL), Livermore, CA (United States), LLNL–TR-2006407, Aug. 2025. doi: 10.2172/2584723, Z. Zhang *et al.*, "FedCSpc: A Cross-Silo Federated Learning System with Error-Bounded Lossy Parameter Compression," *IEEE Transactions on Parallel and Distributed Systems*, vol. 36, no. 7, Jul. 2025, doi: 10.1109/tpds.2025.3564736, K. Purcell, "Using the power of secure Generative AI to eliminate data silos," SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA (United States), Apr. 2024. doi: 10.2172/2345015, A. Miguel Cruz, S. Marshall, C. Daum, H. Perez, J. Hirdes, and L. Liu, "Data silos undermine efforts to characterize, predict, and mitigate dementia-related missing person incidents," *Healthcare Management Forum*, vol. 35, no. 6, pp. 333–338, Nov. 2022, doi: 10.1177/08404704221106156, N. M. Ivanova, A. S. Kashin, and V. P. Ananikov, "Lost Data in Electron Microscopy." arXiv, Aug. 2025. doi: 10.48550/arXiv.2508.18217. Available: https://arxiv.org/abs/2508.18217. [Accessed: Aug. 29, 2025]

[4] S. Sleep, P. Gala, and D. E. Harrison, "Removing silos to enable data-driven decisions: The importance of marketing and IT knowledge, cooperation, and information quality," *Journal of Business Research*, vol. 156, p. 113471, Feb. 2023, doi: 10.1016/j.jbusres.2022.113471

[5] E. Orders, "Stopping Waste, Fraud, and Abuse by Eliminating Information Silos," *The White House*. https://www.whitehouse.gov/presidential-actions/2025/03/stopping-waste-fraud-and-abuse-by-eliminating-information-silos/, Mar. 2025

[6] N. B. Thylstrup, M. Archer, and H. Steiner, "Desiloization and its discontents: The politics of data storage in the age of platformization," *Information, Communication & Society*, vol. 27, no. 13, pp. 2419–2437, Oct. 2024, doi: 10.1080/1369118X.2024.2371803

[7] Particularly, sensitive information, such as security system diagrams, can improve the competitive advantage of the adversary if released unscrupulously.

storage cost, no matter the reason for its inaccessibility. The alternative (data accessible to appropriate end users and stakeholders) has been used to great effect in both the public[1] and private[2] sectors and may even increase if organizations' motivation for eliminating silos becomes value creation, and not legitimacy.[3]

B.3 Best Practice – DMSR Metrics

A consensus exists that quantitative measurement of desired outcomes leads to improved awareness and outcomes. Several entities have advocated for the reporting of metrics related to data quality and availability around EM. Virginia Innovation Partnership Corporation advocated for status indicators regarding the presence and quality of data elements from differing jurisdictions;[4] Yin *et al.* advocated for the automated measurement of information retrieval performance with respect to EM.[5] These proposals illustrate well-crafted metrics with respect to maximal expected utility: they attempt to directly measure the usefulness of such information *to the data user*. More holistic measures, and analysis of such measures in a metrological sense, will improve the overall DMSR ecosystem. This will become especially important as the amount of data and automatically generated content increases, which likely will not coincide with an increase in the amount of useful information available (the so-called "AI Slop" issue[6]).

B.4 Best Practice and Lesson Learned – Curation, Standards, and Interfaces

The naive application of lessons from the previous sections could lead to allowing open access to all data and creating standards and benchmarks to measure the quality and availability of that data. This has been largely enacted in the world of scientific data and has been instrumental in many scientific accomplishments in the development of data science.[7] Several second-order detriments have become apparent in that movement. The *curation* of included data within these open datasets and the specialization of solutions to only these datasets may have a detrimental effect on applicability to real-world situations, the so-called "generalization problem," where

---

[1] P. Huston, V. Edge, and E. Bernier, "Reaping the benefits of Open Data in public health," *Canada Communicable Disease Report*, vol. 45, no. 11, pp. 252–256, Oct. 2019, doi: 10.14745/ccdr.v45i10a01

[2] "Business case for open data resources.data.gov." https://resources.data.gov/resources/open-data/

[3] S. Temiz, M. Holgersson, J. Björkdahl, and M. W. Wallin, "Open data: Lost opportunity or unrealized potential?" *Technovation*, vol. 114, p. 102535, Jun. 2022, doi: 10.1016/j.technovation.2022.102535, J. Borycz *et al.*, "Perceived benefits of open data are improving but scientists still lack resources, skills, and rewards," *Humanities and Social Sciences Communications*, vol. 10, no. 1, p. 339, Jun. 2023, doi: 10.1057/s41599-023-01831-7

[4] Dan Cotter, "The Community Lifeline Status System: The Foundation for the EOC of the Future," presentation, June 25, 2024.

[5] K. Yin *et al.*, "DisastIR: A Comprehensive Information Retrieval Benchmark for Disaster Management." arXiv, Sep. 2025. doi: 10.48550/arXiv.2505.15856. Available: https://arxiv.org/abs/2505.15856. [Accessed: Nov. 01, 2025]

[6] S. Paech, A. Roush, J. Goldfeder, and R. Shwartz-Ziv, "Antislop: A Comprehensive Framework for Identifying and Eliminating Repetitive Patterns in Language Models." arXiv, Oct. 2025. doi: 10.48550/arXiv.2510.15061. Available: https://arxiv.org/abs/2510.15061. [Accessed: Nov. 18, 2025], C. Shaib, T. Chakrabarty, D. Garcia-Olano, and B. C. Wallace, "Measuring AI "Slop" in Text." arXiv, Sep. 2025. doi: 10.48550/arXiv.2509.19163. Available: https://arxiv.org/abs/2509.19163. [Accessed: Nov. 18, 2025]

[7] I. Peters, P. Kraker, E. Lex, C. Gumpenberger, and J. I. Gorraiz, "Zenodo in the Spotlight of Traditional and New Metrics," *Frontiers in Research Metrics and Analytics*, vol. 2, Dec. 2017, doi: 10.3389/frma.2017.00013, I. D. R. Crespo Garrido, M. Loureiro García, and J. Gutleber, "The Value of an Open Scientific Data and Documentation Platform in a Global Project: The Case of Zenodo," in *The Economics of Big Science 2.0*, J. Gutleber and P. Charitos, Eds., Cham: Springer Nature Switzerland, 2025, pp. 181–200. doi: 10.1007/978-3-031-60931-2_14, M.-A. Sicilia, E. García-Barriocanal, and S. Sánchez-Alonso, "Community Curation in Open Dataset Repositories: Insights from Zenodo," *Procedia Computer Science*, vol. 106, pp. 54–60, Jan. 2017, doi: 10.1016/j.procs.2017.03.009

specialization to a dataset reduces performance in realistic situations.[1] The curation of the information kept about each datum by keeping only "features" of the data which are requested by data standards is another understandable application of the frameworks above. Again, the specialization of these "features" to the needs and standards of the present has a detrimental effect; decades of empirical evidence show the dominance of "raw" data over feature-curated data in fields from chemical modeling to computer vision, especially in fields with fast-moving question cycles like forensics.[2] Furthermore, the era of LLMs has drastically reduced the effort required to retrieve data from large, natural language corpuses, making "feature" curation largely obsolete. Both lessons learned illustrate what happens when the ultimate mission goal of DMSR—to have utility for future applications—is subverted by shortcuts to improving other metrics.

Instead, the concept of interfaces can provide a better application of the frameworks and best practices already laid out. Interfaces are lightweight applications that enable standardized communication patterns between the data requester and the data storer; the latter can store the data in ways most convenient or useful for their purposes without impacting the data usability for the former. These interfaces have two critical differences from data standards. Standards disincentivize the data storer from keeping the original data, whereas interfaces are agnostic to that question—it is likely most convenient for the storer to keep it in its original form. Interfaces are also dynamic compared to standards. Updating a standard either requires deprecating all data before the standard change or migrating data to the new standard. An interface is simply an application that transforms data on request and, therefore, can be updated with very low effort. Software compilers are an excellent example of interfaces, where computing hardware manufacturers and software developers can improve their own product independently, while still being confident that it will work because of the defined interfaces. Virginia Innovation Partnership Corporation is an excellent example of advocacy for data interfaces in the EM sector.

[1] Y. Y. Bay and K. A. Yearick, "Machine Learning vs Deep Learning: The Generalization Problem." arXiv, Mar. 2024. doi: 10.48550/arXiv.2403.01621. Available: https://arxiv.org/abs/2403.01621. [Accessed: Nov. 13, 2025], J. Fodor, "Line Goes Up? Inherent Limitations of Benchmarks for Evaluating Large Language Models." arXiv, Feb. 2025. doi: 10.48550/arXiv.2502.14318. Available: https://arxiv.org/abs/2502.14318. [Accessed: Nov. 13, 2025]

[2] M. Girard* *et al.*, "Uranium Oxide Synthetic Pathway Discernment through Unsupervised Morphological Analysis," *Journal of Nuclear Material*, 2021. A. Hagen, "Training and using neural networks in flexible material analysis workflows for nuclear forensics research." Ames, Iowa, Sep. 2025.

Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99354

1-888-375-PNNL (7665)

www.pnnl.gov