RESEARCH REPORT

AI and the Future of Emergency Management

Market Supply and Adoption Pathways

JESSICA JENSEN, JESSIE RIPOSO, LEAH DION, GLEN L WOODBURY



This research was funded through a generous grant from

AIDE A.I. for Disasters + Emergencies
The Markle Foundation

For more information on this publication, visit **www.rand.org/t/RRA4625-1**.

About RAND

RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our research integrity is grounded in RAND's core values of quality and objectivity. Rigorous quality assurance procedures, conflict of interest screening, and transparency in funding ensure that every study is objective and nonpartisan. Learn more at www.rand.org/integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.
© 2026 RAND Corporation
RAND® is a registered trademark.

Cover: Image--witsarut sakorn/Getty Images, concentric lines--arbaz bagwan/Getty Images.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, visit www.rand.org/about/publishing/permissions.

RR-A4625-1

About This Report

In this report, we present the results of a landscape assessment of artificial intelligence (AI)– enabled products that are available to support the work of emergency management (EM) across the United States. The Markle Foundation commissioned this assessment as part of its AI for Disasters and Emergencies (AIDE) Initiative, a multi-sector, bipartisan collaboration between the Foundation, Aspen Digital (a program of the Aspen Institute), and RAND. We identified 1,179 AI-enabled products that span preparedness, response, and recovery and assessed this landscape against the operational and adoption conditions that matter for state, local, tribal, and territorial governments and the nonprofit organizations that work alongside them. We characterize the products we identified, describe what it takes to adopt those products, and identify what stands between the availability of those products and their adoption. We do not rank or evaluate specific products. This report is intended to give the EM community, technology companies, funders, policymakers, and researchers a reliable picture of available AI-enabled products so that they can act on what the landscape suggests for their own decisions, investments, and further work.

Homeland Security Research Division

This work was conducted within the RAND Homeland Security Research Division (HSRD). For more information on HSRD, visit www.rand.org/hsrd.

For more information on this publication, see www.rand.org/t/RRA4625-1.

Funding

This research was funded through a generous grant from the Markle Foundation.

Acknowledgements

We thank the Markle Foundation for sponsoring the AIDE Initiative and for funding a first-of-its-kind landscape assessment of AI-enabled products to support the work of EM. The challenges that are inherent in assessing a vast and rapidly moving product market at scale are real, and we are grateful to the Foundation for recognizing how critical such an assessment is to thinking about the potential applications of AI in EM and its capacity to evolve the field.

We are grateful to William Marcellino, Joie Acosta, Patrick Roberts, Andrew Lauland, and Jason Thomas Barnosky for their guidance in the formative stages of the project. Prateek Puri provided thought leadership and developed and executed the code that supported web searches and commercial company data application programming interface (API) queries. Shawn Chao and Mohammad Ahmadi did thoughtful, quality-driven, and persistent work on the API code used to collect product data and on the automated verification processes that made that collection reliable. James Caridi-Doyle wrote the code that cleaned our data after each collection. Shawn Chao, Mohammad Ahmadi, James Caridi-Doyle, Rachel Slover, Nathaniel Tunggal, and Quintessa Guengerich devoted substantial time to supporting the human verification of the data—from training through assessment and iterative discussion—and several verifiers returned for more than one round. Quintessa Guengerich also conducted a careful analysis of the results of our human verification processes that gave us the confidence to rely on the data we had gathered.

We thank the full team for their early and ongoing efforts to develop a reliable approach to estimating product costs, and we are especially grateful to Thomas Goughnour and Bradley Wilson for tackling

the cost estimation work, exploring multiple approaches, and helping us think through how to identify promising products.

We thank representatives of the AIDE Initiative from the Markle Foundation and Aspen Digital for their support and insights throughout this project. We also thank our peer reviewers: Patrick Roberts, Peter Schirmer, and Giovanni Malloy; Jason Thomas Barnosky, associate director of the Infrastructure, Immigration, and Security Operations Program at RAND, who reviewed the report in his official program capacity; Eric Landree, acting vice president of HSRD; and Jeremy Eckhause, research quality assurance manager within HSRD. The thoughtful feedback from all these reviewers improved the report, and we are grateful for it.

Summary

Issue

Many publications describe what artificial intelligence (AI) could do for emergency management (EM) someday. Almost none of those publications describe the AI-enabled products that are available or what it would take to adopt those products. This information gap matters: Emergency managers and others who support EM functions are making procurement decisions in a crowded technology marketplace with limited time, technical capacity, and standardized information about how AI-enabled products perform against the realities of their work. Without a clearer understanding of the market, they risk spending scarce funds on products they cannot integrate, sustain, or rely on. They may also overlook products that could meaningfully extend limited resource capacities, improve disaster response outcomes, accelerate recovery, and improve service delivery to affected communities. As hazards and disasters increase in intensity and frequency, the demands on the resources of emergency managers and others who support EM functions often outpace their capabilities and capacity. The Markle Foundation commissioned this landscape assessment as part of its AI for Disasters and Emergencies (AIDE) Initiative to give those decisionmakers, along with technology companies, funders, policymakers, and researchers, a picture of the AI-enabled product landscape and the conditions that will shape how those products can support EM.

Approach

We used a mixed-methods approach to collect and analyze product data between October 2025 and March 2026. We first developed a framework of 45 EM task areas that span preparedness, response, and recovery, drawn from authoritative sources and validated through expert review and engagement with emergency managers. We then identified products through complementary search strategies, including the iterative use of large language model (LLM) application programming interfaces (APIs), commercial company data APIs, structured web searches, conference exhibitor lists, and stakeholder recommendations. The initial search surfaced 1,892 candidate products, which we narrowed through deduplication and scoping filters to 1,179 products from 717 technology companies. We developed a structured framework of product features that could be searched for consistently from public sources and collected information using a multi-agent LLM pipeline paired with a multistage verification process that combined automated tools, expert review, and formal RAND quality assurance. Finally, we applied descriptive statistics and qualitative analysis to characterize the product landscape, identify patterns, and assess the conditions for product adoption and diffusion across EM.

Key Findings

The AI-enabled product landscape is substantial but unevenly distributed among EM functional areas. We identified 1,179 AI-enabled products with potential relevance to EM. The market heavily favors response (53 percent of products), moderately serves recovery (36 percent of products), and thinly covers preparedness (12 percent of products). Approximately half of the products we identified are either purpose-built for EM or have documented use in EM contexts; more than 40 percent are general-purpose solutions with no EM-specific design. All 45 EM task areas have at least nine associated products, but coverage within the task areas is partial: Most products address narrow slices of work and would require stacking with other products to meet the full scope of a task area. Clear gaps remain: Multi-organization coordination, preparedness work,

and survivor-facing services are consistently underserved. These gaps create meaningful opportunities for technology companies.

Procurement realities create meaningful distance between availability and adoption; concerted efforts by the user community and technology providers could close it. One in five products depends on another product to function. Most products require technical integration that presupposes information technology (IT) capacity, ongoing data inputs, and continuous internet connectivity, all of which disasters can disrupt. Pricing is largely opaque: Most technology companies do not publish prices, and assembling reliable cost information across the product landscape was not feasible even with substantial research effort. Sixty-eight percent of the technology companies associated with the products identified in our analysis do not publicly advertise 24/7 technical assistance. Privacy risk is near-universal across the landscape; legal and cyber concerns affect a meaningful minority of products, and verifiable AI governance assurance is rare.

Adoption and diffusion conditions for AI in EM are in an early emergence stage, which is characteristic of technology domains in which innovation outpaces governance and policy. The available evidence suggests that most state, local, tribal, and territorial (SLTT) EM offices remain in the earliest stages of AI adoption and that adoption is concentrated in larger and better-resourced jurisdictions. Funding, staffing, IT capacity, and procurement infrastructure are the most persistent constraints, and they echo barriers that have been documented across decades of EM research. Adoption across organizations that are adjacent to EM, such as utilities, nonprofits, and public health agencies, appears further along than in EM offices but is uneven and concentrated in administrative and communication uses. Infrastructure that supports the broad diffusion of new technology in a field, such as demonstration, peer learning, institutional norms, and state-to-local provisioning, is limited for AI. The opportunity to move toward the broader adoption and diffusion of AI-enabled products exists, but realizing it will require concerted action.

Limitations

Our assessment characterizes the AI-enabled products that surfaced from our search strategies and not from the universe of products that could exist in the EM space. Our analysis is based almost entirely on publicly available information; characterizations requiring nonpublic information were not feasible at scale. The use of LLMs in product identification favors products with a strong online presence and clear AI-related language. Some product characterizations are based on inference from available evidence rather than direct technology company confirmation, and information that was current at the time of data collection may already be outdated given how rapidly the market moves. In this report, we characterize what products exist and the conditions affecting their adoption readiness; we do not directly observe what products are being used in specific SLTT jurisdictions or nonprofit organizations or how they perform. Product cost was the dimension in which the limitations in available data were most apparent, and no specific cost figures are reported.

Recommendations and Next Steps

The findings suggest that progress toward the broad adoption and diffusion of AI-enabled products in EM can be achieved through the decisions and purposeful effort of multiple stakeholder groups, each working in the areas in which they have unique influence. The report provides sequenced action agendas for emergency managers and others who support EM functions, technology companies (with a separate agenda for big tech and hyperscalers), policymakers, and researchers. We identify things each stakeholder group can do in the next 12 months, the near term (i.e., the next one to three years), and on an ongoing basis to support the adoption and diffusion of AI-enabled products to support EM. Philanthropic organizations, professional

associations, and other intermediaries are uniquely positioned to support progress across all three horizons and stakeholder groups.

Beyond the sequenced agendas for each stakeholder group, four high-impact initiatives are feasible within the next 12 months and could meaningfully accelerate adoption and diffusion at scale if pursued in close collaboration with technology companies, philanthropy organizations, emergency managers, the broader set of organizations that perform EM functions, and policymakers. The initiatives are as follows:

- Establish a standing cross-sector alliance to focus sustained attention on AI-enabled product adoption and diffusion, anchored by members with the capacity to financially support and sustain the alliance.
- Launch a national-level EM AI readiness and adoption accelerator to function as an independent, Consumer Reports–style clearinghouse and provide product evaluations, standardized buyer guides, model procurement and contracting language, peer-learning cohorts, and practical deployment templates for low-capacity jurisdictions.
- Develop and sustain a national, freely accessible education and training program designed as a tiered learning pathway to build AI awareness and capability among emergency managers and others who perform EM functions.
- Fund the critical outstanding research questions identified in this report—including whether LLMs can effectively support EM tasks and whether national AI governance, procurement, and IT policies create systematic barriers to adoption—so that the answers can inform pilot projects, alliance work, and policy decisions.

Contents

About This Report iii
Summary v
Figures and Tables xi

CHAPTER 1
Introduction 1
EM Encompasses a Broad Set of Tasks Performed by a Variety of People and Organizations 1
Technology, Particularly AI, May Evolve EM Practice and Overcome Some of These Challenges 3
To Maximize the Opportunity That AI-Enabled Products Afford, They Must Be Adopted Within Individual Organizations Across the Country 4
The Markle Foundation Charged RAND with Assessing the Landscape of AI-Enabled Products That Are Available to Support EM 7
We Share the Information That Led to This Conclusion Across Four Additional Chapters and Nine Appendixes 9

CHAPTER 2
Many Products Exist with the Potential to Support Emergency Management 11
Many Products Are Purpose-Built for EM (or Have Been Used to Support It) and Are Hazard Agnostic 11
There Are Nine or More Products Available to Support Every Functional Area and Task Area 13
Most Products Support Multiple Task Areas 13
Products Supporting Task Areas That Are Important to Emergency Managers Exist 15
There Are Products Other Entities That Support EM Functions Would Use 21
Nearly Half of the Products Identified in This Research Are General Solutions That Have the Potential to Support EM Task Areas 24
Many Products Support Specific Tasks Within an EM Task Area but Would Require Stacking to Best Meet EM Needs 27
There Are Clear Gaps in the Product Landscape 28
Some Products Offer Novel Capabilities That Address EM Needs 29
Conclusion 30

CHAPTER 3
Procurement Realities: From Available to Adopted 31
One in Five Products in This Landscape Is a Component That Depends on Another Product to Function 31
Most Products Are Delivered as Cloud Services, and Many Offer Additional Deployment Options 32
The Path from Purchase to a Working, Sustained Product Is Often Neither a Single Transaction nor a One-Time Setup 33
Most Products Stop Working When Connectivity Fails 35
Most Technology Companies Provide Learning Support for Users but Not 24/7 Technical Assistance 36
Pricing Is Largely Opaque, and Access Models Vary Widely 37
Most Products Have Documented Deployment Across Multiple Sites or Countries 38
The Type and Degree of Risk Vary Across Products 40
Conclusion 45

CHAPTER 4
From Awareness to Broad Adoption: What Will Shape AI Diffusion in Emergency Management? 47
Adoption Within EM Offices Is Emerging at Different Rates 47
Diffusion Is Possible in EM, but the Conditions That Enable It Are Not Available Yet for AI 49
AI Adoption Across Organizations That Support EM Is Meaningful but Emerging at Different Rates 51
Technology Company Design Choices and Communication Practices Shape What Practitioners Can Do 53
Institutional Conditions Shape What Practitioners Can Realistically Do 55
Conclusion 57

CHAPTER 5
From Awareness to Broad Adoption and from Adoption to Diffusion: An Agenda for Action 59
A Sequenced Agenda for Action 59
High-Impact Initiatives to Accelerate Adoption and Diffusion at Scale 60
Conclusion 62

APPENDIXES
A. A Brief History of AI 63
B. Project Approach 67
C. Response Task Areas and AI-Related Patterns 79
D. Disaster Recovery Task Areas and AI-Related Patterns 87
E. Preparedness Task Areas and AI-Related Patterns 95
F. Examples of Observed Response, Recovery, and Preparedness Product Gaps 103
G. Sequenced Action Agendas for Stakeholder Groups to Support Adoption and Diffusion 107
H. Criteria Practitioners Can Use to Evaluate Their Organization's Adoption Readiness Based on Product Characteristics 111
I. Questions for Product Technology Companies: A Checklist for Practitioners 117

Abbreviations 129
References 131

Figures and Tables

Figures

1.1. AI Subfields 4
1.2. Organizational Technology Adoption 6
1.3. Field Level Technology Diffusion 6
2.1. The Extent to Which Products Are Purpose-Built for EM Overall and by Functional Area 12
2.2. Task Areas with the Highest and Lowest Number of Products by Functional Area 14
2.3. Number of Task Areas Each Product Supports 15
3.1. Product Burden 34
3.2. Potential Connectivity Risk 36
3.3. Provision of Technical Assistance and Learning Support 37
3.4. Product Deployment Scale 39
3.5. Potential Risk from Evidence of Privacy, Legal, or Cyber Concerns 40
4.1. Local and State EM AI Access 48
B.1. Workflow for Product Identification, Information Extraction, Validation, and Analytic Dataset Preparation 67
H.1. Questions About Readiness to Procure and Adopt AI-Enabled Products 112
I.1. Questions About the Applicability of AI-Enabled Products 118

Tables

2.1. Situational Awareness, Hazard Monitoring, and Information Sharing: Product Clusters, What Products Do, and Example Products 16
2.2. Incident Coordination and Operational Decision Support: Product Clusters, What Products Do, and Example Products 17
2.3. Damage Assessment: Product Clusters, What Products Do, and Example Products 18
2.4. Public Information, Alerts, and Protective Guidance: Product Clusters, What Products Do, and Example Products 19
2.5. Protective Action Decision and Implementation: Product Clusters, What Products Do, and Example Products 20
2.6. Evacuation and Transportation Operations: Product Clusters, What Products Do, and Example Products 21
2.7. Critical Infrastructure and Essential Services Stabilization: Product Clusters, What Products Do, and Example Products 22
2.8. Volunteer and Donations Management: Product Clusters, What Products Do, and Example Products 23
2.9. General Tasks That AI-Enabled Products Might Support, Product Type Descriptions, and Example Products 26
3.1. Examples of User-Required Data Inputs and Associated Product Types 35
3.2. Most Observed Patterns in the Privacy Concerns Data 41
3.3. Most-Observed Patterns in the Potential Legal Concerns Data 42
3.4. Most-Observed Patterns in the Potential Cyber Concerns Data 43
C.1. Response Task Areas and AI-Related Patterns 80
D.1. Disaster Recovery Task Areas and AI-Related Patterns 88
E.1. Preparedness Task Areas and Patterns Regarding What Current AI Products Enable 96
F.1. Examples of Observed Response, Recovery, and Preparedness Product Gaps 103

G.1. Sequenced Agenda for Emergency Managers and Others Who Support EM Functions 107
G.2. Sequenced Agenda for Technology Companies 108
G.3. Sequenced Agenda for Big Tech and Hyperscalers 109
G.4. Sequenced Agenda for Policymakers 110
G.5. Sequenced Agenda for Researchers 110

CHAPTER 1

Introduction

Many publications describe what artificial intelligence (AI) could do for emergency management (EM) someday.[1] Almost none of those publications describe the AI-enabled products that are available or what it would take to adopt those products. This is because in a fast-moving marketplace, product features rapidly change, and companies pivot, get acquired, or sunset products. These realities make it very challenging to identify and study such products. In this report, we present the results of an effort to overcome those challenges and address this information gap. Based on the best information we could find, we document what AI-enabled products are available, what it takes to adopt those products, what stands in the way of their adoption, and potential steps to advance their adoption and diffusion.

Emergency managers and others who support EM functions are making procurement decisions in a crowded technology company landscape with limited time, technical capacity, and standardized information about how AI-enabled products perform against the realities of their work. In this report, we do not resolve those decisions at the level of individual products; the landscape is too broad for that. Instead, we characterize this landscape as a whole; assess it against the operational and adoption conditions that matter for state, local, tribal, and territorial (SLTT) governments and the nonprofit organizations that work alongside them; and identify in which EM functional and task areas products exist, where they do not, and what may stand between product availability and adoption at both the organizational level and the level of broad field diffusion. We also offer practical next steps for the public, nonprofit, and private sectors.

In the remainder of this chapter, we establish what a reader needs to understand our assessment: what EM is and who carries it out, the conditions under which EM must currently operate, what we mean by *AI* and why its potential matters, what the adoption of AI-enabled products actually requires, and the remit from the Markle Foundation for this assessment.

EM Encompasses a Broad Set of Tasks Performed by a Variety of People and Organizations

The personnel who work in SLTT EM offices are "charged with creating the framework within which [organizations and] communities reduce vulnerability to hazards and cope with disasters."[2] The mission of these offices and personnel is to "coordinat[e] and integrat[e] all activities necessary to build, sustain, and improve

[1] For example, see, Federal Emergency Management Agency (FEMA), *Strategic Foresight 2050: Final Report*, November 2024; Xiao Liang, "Image-Based Post-Disaster Inspection of Reinforced Concrete Bridge Systems Using Deep Learning with Bayesian Optimization," *Computer-Aided Civil and Infrastructure Engineering*, Vol. 35, No. 5, May 2019; U.S. Department of Homeland Security, *Department of Homeland Security Artificial Intelligence Roadmap 2024*, 2024; Jonathan Barr, Alexander Hagen, Grant Tietje, Nicholas Betzsold, Alex Greer, and Eric Best, *Emergency Management of Tomorrow Research: Artificial Intelligence Landscape Assessment*, Science and Technology Directorate, U.S. Department of Homeland Security, and Pacific Northwest National Laboratory, June 2024.

[2] FEMA, "Principles of Emergency Management," 2007.

the capability to mitigate against, prepare for, respond to, and recover from threatened or actual natural [hazards], acts of terrorism, or other man-made disasters."[3]

This mission is not carried out by any single office or set of professionals. The work of EM is, at its core, a distributed function—one in which people who are employed in EM positions facilitate frameworks, but the actual work of preparedness, response, recovery, and mitigation is carried out by a variety of entities.[4]

Public health agencies, school systems, hospitals, utility operators, voluntary organizations, faith-based groups, law enforcement, public works departments, and many others have essential roles. They all bring expertise, resources, and authority, without which the EM mission cannot be accomplished.[5]

The people in EM offices and positions, whom we refer to as *emergency managers*, serve as connectors and coordinators within this broader system. They facilitate hazard mitigation planning, help jurisdictions and organizations prepare through training and exercises, and coordinate response and recovery across many actors when hazard events occur. EM is what happens across the entire distributed system, not only in EM offices.

The work that supports EM is broad in scope and complex in execution. Such work spans the EM functional areas of preparedness, response, and recovery,[6] and each functional area is made up of multiple task areas that must be carried out at a jurisdictional level when hazard events are anticipated, ongoing, or past (see the box). We identified 45 broad EM task area categories: 20 response task areas, 16 recovery task areas, and nine preparedness task areas. We provide descriptions of these task areas in Appendixes C through E.

[3] FEMA, 2007.

EM Functional Areas

Disaster preparedness "refers to the actions taken to plan, organize, equip, train, and exercise to build and sustain the capabilities necessary to prevent, protect against, mitigate the effects of, respond to, and recover from those threats that pose the greatest risk to the security of the Nation."[a]

Disaster response refers to actions taken when a hazard event is imminent or occurring to save people, property, and the environment. The response does not end until the immediate threat to people, property, and the environment has ceased and an initial damage assessment is complete.[b]

Disaster recovery "defines capabilities necessary for communities affected or threatened by any incident to rebuild infrastructure systems, provide adequate, accessible interim and long-term housing that meets the needs of all survivors, revitalize health systems (including behavioral health) and social and community services, promote economic development, and restore natural and cultural resources."[c]

[a] Barack Obama, Presidential Policy Directive 8: National Preparedness, March 30, 2011.
[b] FEMA, *National Preparedness Goal*, 2nd ed., U.S. Department of Homeland Security, September 2015; FEMA, *National Response Framework*, 4th ed., U.S. Department of Homeland Security, October 28, 2019.
[c] FEMA, *National Disaster Recovery Framework*, 2nd ed., U.S. Department of Homeland Security, June 2016, p. 4.

[4] Jessica Jensen, Sarah Bundy, Brian Thomas, and Mariama Yakubu, "The County Emergency Manager's Role in Disaster Recovery," *International Journal of Mass Emergencies and Disasters*, Vol. 32, No. 1, March 2014.

[5] Jensen et al., 2014.

[6] *Mitigation*, which overlaps with but is distinct from the EM functional areas that this assessment covers, was intentionally placed outside the scope.

All together, these task areas illustrate a scope of work that any single jurisdiction is nominally responsible for and that, in practice, no single office has the capacity to carry out alone.

The Demands for EM Are Outpacing Capacity

EM has long been under-resourced, particularly at SLTT levels, where the primary responsibility for managing hazards, vulnerabilities, and the consequences of hazard events resides. In most jurisdictions, the people, funds, or other resources needed to meet related EM demands have been insufficient.[7] Now, the demands for EM are outpacing capacity. EM offices are being asked more frequently to take on activities that have not historically been part of their roles,[8] and they must prepare for an increasing number of hazard types.[9] Hazard events are increasing in number and severity:[10] Events are cascading and increasingly converging into polycrises,[11] adding new complexity to an already challenging operating environment.[12]

At the same time, the federal government's current policy is to shift more of the financial and managerial burden of EM to SLTT levels.[13] These shifts are still underway, and the ultimate funding picture is uncertain. All these changes also affect the organizations that support EM functions.

Technology, Particularly AI, May Evolve EM Practice and Overcome Some of These Challenges

AI-enabled products are already augmenting a variety of EM functions and could increasingly do so. For the purposes of this report, *AI refers to technology enabled by computational systems that are capable of performing one or more specific tasks within one or more task areas that would require intelligence if carried out by humans*—such as reasoning, prediction, problem-solving, learning, planning, perception, and communication. AI draws on computer science, mathematics, statistics, engineering, cognitive science, and related disciplines to develop algorithms and architectures that mimic or augment human cognitive functions. Although often discussed as a recent development, AI as a concept and in practice has been around for decades (see

[7] National Preparedness Analytics Center, *Emergency Management Organizational Structures, Staffing, and Capacity Study: State, Local, and Territory Findings Report*, Argonne National Laboratory, July 2025; John R. Labadie, "Problems in Local Emergency Management," *Environmental Management*, Vol. 8, No. 6, November 1984; William J. Petak, "Emergency Management: A Challenge for Public Administration," *Public Administration Review*, Vol. 45, January 1985; Chris Currie, *Disaster Assistance High-Risk Series: State and Local Response Capabilities*, U.S. Government Accountability Office, GAO-26-108599, December 18, 2025.

[8] National Preparedness Analytics Center, 2025; Natural Hazards Center, "State of Emergency—Resource Commitments and Compassion in an Age of Extremes," video, July 13, 2023.

[9] For a discussion of AI as an emergent risk, see Dan Hendrycks, Mantas Mazeika, and Thomas Woodside, "An Overview of Catastrophic AI Risks," arXiv, arXiv:2306.12001v6, October 9, 2023.

[10] Currie, 2025.

[11] *Polycrisis* refers to a situation in which multiple crises occur simultaneously and interact with one another, such that their combined effects are more severe, more complex, and harder to manage than each crisis in isolation.

[12] World Economic Forum, *The Global Risks Report 2023*, 18th ed., January 2023.

[13] Page Forrest and Jad Maayah, "Uncertainty Surrounding Federal Disaster Funding Looms over State Budgets," Pew Charitable Trusts, January 15, 2026; Sara McTarnaghan, Lucy Dadayan, and Andrew Rumbach, "The Trump Administration Wants to Shrink the Federal Government's Role in Disaster Management. Do States Have the Fiscal Capacity to Weather the Storm?" Urban Institute, July 15, 2025.

FIGURE 1.1
AI Subfields



NOTE: GPT = generative pretrained transformer; LLM = large language model; NLP = natural language processing.

Appendix A) and has evolved to include several subfields that overlap somewhat (Figure 1.1).[14] The names of these subfields are often used interchangeably with the term AI. Within each of these subfields, methods and techniques have been developed to enable computational systems to reason, learn, solve, or complete other tasks requiring intelligence.[15]

An AI-enabled product is a product that performs a task that requires intelligence by using one or more of the techniques encompassed by these subfields. AI-enabled products matter in EM because they have the potential to improve the decisions and actions taken by emergency managers and others who are responsible for preparing for, responding to, and recovering from disasters. Before disasters occur, these products can reduce the amount of time personnel spend on routine, desk-bound tasks and enable more time to be spent on engaging with communities, strengthening preparedness, and advocating for mitigation. During disasters, the products can provide more-timely, higher-quality, and better-integrated information to support operational decisionmaking and action. After disasters, the products can support the execution of recovery activities in ways that are more efficient, comprehensive, and just. The products would not produce these outcomes—they would support the people who do produce these outcomes. AI-enabled products can extend the capacity of EM personnel and others who support EM functions to operate more efficiently and effectively as product adoption and diffusion occurs.

To Maximize the Opportunity That AI-Enabled Products Afford, They Must Be Adopted Within Individual Organizations Across the Country

For AI to meaningfully strengthen EM, two things must happen. First, individual organizations must move new products from initial interest into sustained use (Figure 1.2).[16] Second, whole fields must reach a state in which adoption is widespread enough that a technology becomes part of how the work is done, not an exception to how the work is done. Neither of these things can happen quickly and neither can happen automatically. Both depend on conditions that can be supported or obstructed by factors inside and outside any single organization.

In practice, this progression is iterative and nonlinear. Sometimes organizations fail to progress beyond initial implementation to sustained use or full adoption. Across a field, technology diffusion is associated with the emergence of several conditions rather than a linear process (Figure 1.3).[17]

---

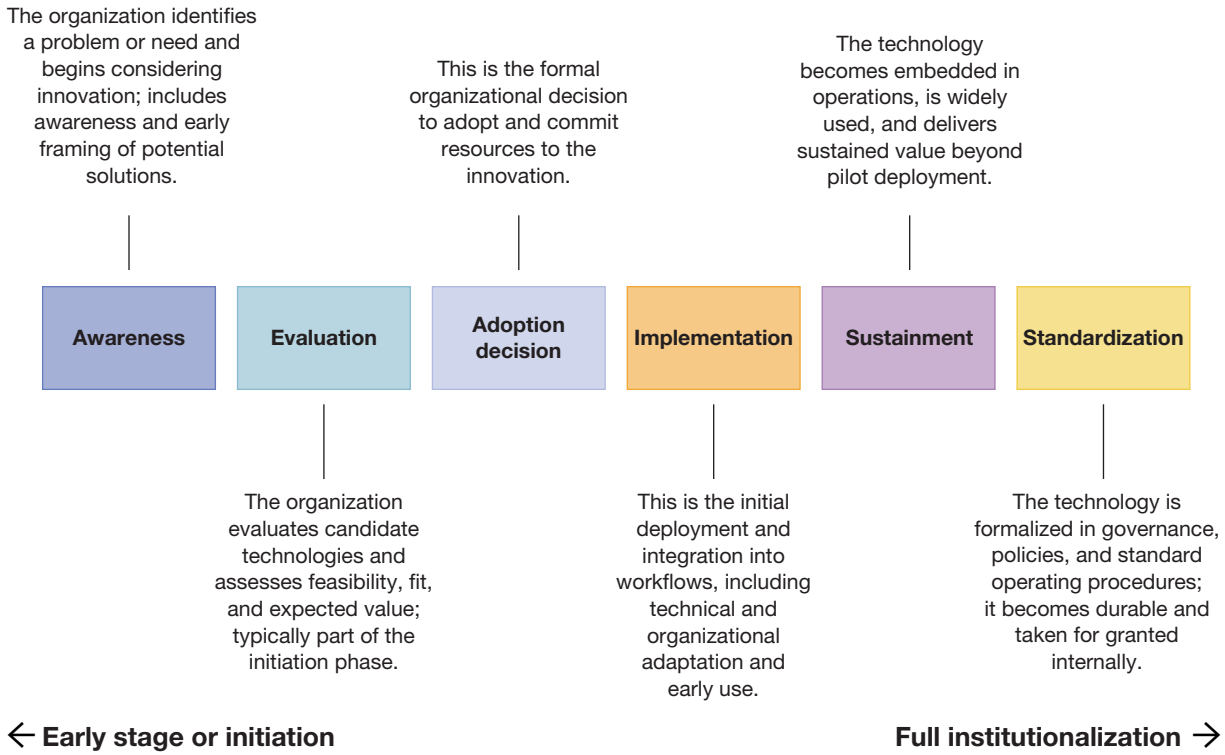
[14] Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed., Pearson, 2022.

[15] Stephen Lucci, Sarhan M. Musa, and Danny Kopec, *Artificial Intelligence in the 21st Century: A Living Introduction*, 3rd ed., Mercury Learning and Information, 2022.

[16] Robert G. Fichman and Chris F. Kemerer, "The Assimilation of Software Process Innovations: An Organizational Learning Perspective," *Management Science*, Vol. 43, No. 10, October 1997; Fariborz Damanpour and Marguerite Schneider, "Phases of the Adoption of Innovation in Organizations: Effects of Environment, Organization and Top Managers," *British Journal of Management*, Vol. 17, No. 3, September 2006; Michael J. Gallivan, "Organizational Adoption and Assimilation of Complex Technological Innovations: Development and Application of a New Framework," *Data Base for Advances in Information Systems*, Vol. 32, No. 3, Summer 2001; Everett M. Rogers, *Diffusion of Innovations*, 5th ed., Free Press, 2003; Pamela S. Tolbert and Lynne G. Zucker, "Institutional Sources of Change in the Formal Structure of Organizations: The Diffusion of Civil Service Reform, 1880–1935," *Administrative Science Quarterly*, Vol. 28, No. 1, March 1983; Kevin Zhu, Kenneth L. Kraemer, and Sean Xu, "The Process of Innovation Assimilation by Firms in Different Countries: A Technology Diffusion Perspective on E-Business," *Management Science*, Vol. 52, No. 10, October 2006.

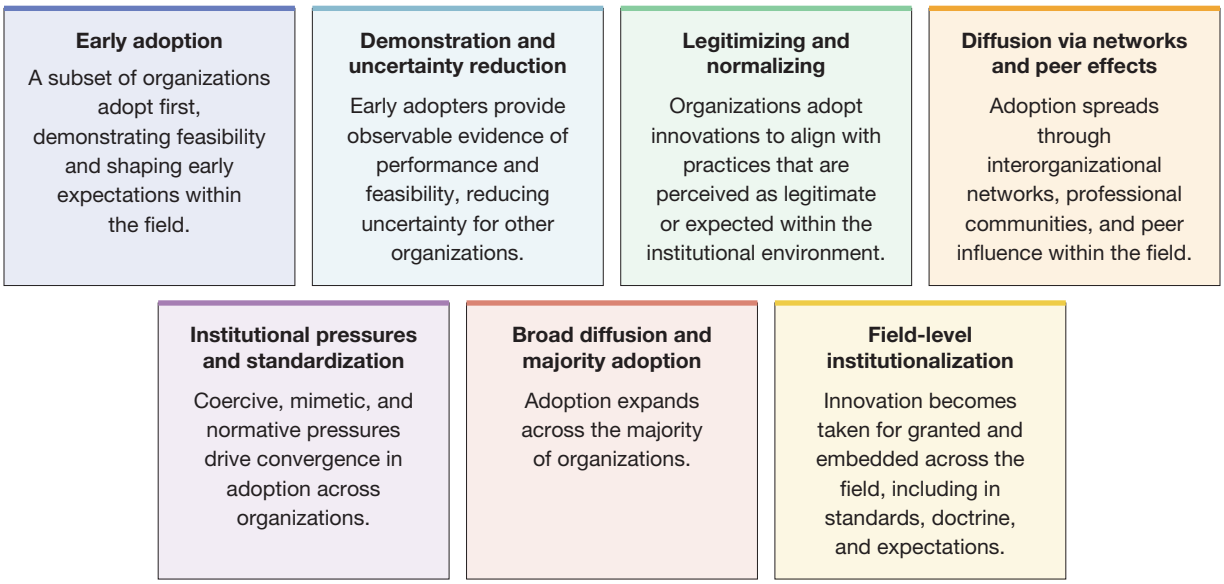
[17] See Paul J. DiMaggio and Walter W. Powell, "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields," *American Sociological Review*, Vol. 48, No. 2, April 1983; Trisha Greenhalgh, Glenn Robert, Fraser Macfarlane, Paul Bate, and Olivia Kyriakidou, "Diffusion of Innovations in Service Organizations: Systematic Review and Recommendations," *Milbank Quarterly*, Vol. 82, No. 4, December 2004; John W. Meyer and Brian Rowan, "Institutional-

FIGURE 1.2
Organizational Technology Adoption



NOTE: Each stage represents a qualitatively distinct phase of organizational adoption. Stages may overlap in practice.

FIGURE 1.3
Field Level Technology Diffusion



The relative importance of these mechanisms varies across contexts. In some fields, diffusion is driven primarily by demonstrated performance and peer learning. In others, mandates or normative expectations accelerate adoption. In still others, adoption lags because the conditions that support any of these mechanisms are weak. The use of low-risk pilots can bridge gaps between vendor claims and performance reality during implementation and sustainment, although the criteria for success should be clearly structured. The mechanisms are best understood not as required steps in a fixed sequence but as indicators and drivers of progression toward broader field-level adoption.

The intraorganizational process and field-level diffusion are linked. Without individual organizations successfully integrating technology into sustained use, there is no evidence for peers to observe and nothing for norms or standards to coalesce around. Without field-level diffusion, individual organizational adoptions remain isolated wins that do not change the broader practice of EM.

The Markle Foundation Charged RAND with Assessing the Landscape of AI-Enabled Products That Are Available to Support EM

The Markle Foundation, founded in 1927, works "to improve the economic well-being, health, and national security of all Americans through catalyzing cross-sector, bipartisan solutions that address critical national priorities."[18] Since 1998, the Foundation's work has focused on realizing the potential of information technology to address public problems.[19]

In 2025, the Foundation launched the AI for Disasters and Emergencies (AIDE) Initiative, which is described on its homepage as "an effort to leverage AI for good to help emergency managers and communities better prepare for, respond to, and recover from natural disasters."[20] AIDE is a "multi-sector, nonpartisan" collaboration between the Foundation, Aspen Digital (a program of the Aspen Institute), and RAND.[21] AIDE was created to identify opportunities to support the broad adoption and diffusion of AI-enabled products in EM. In support of AIDE, the Foundation charged RAND with conducting a landscape assessment of AI-enabled products that are available to help EM.[22] The Foundation asked RAND to focus on three functional areas (i.e., preparedness, response, and recovery) and requested that we cast a broad net within those areas, so as not to exclude technological solutions that might support EM tasks even if those solutions are not built for EM specifically. The Foundation asked for the fullest possible picture of what the AI product market looks like today rather than focusing narrowly on products to support certain hazards, task areas, or product categories where AI attention could be visibly concentrated.[23]

ized Organizations: Formal Structure as Myth and Ceremony," *American Journal of Sociology*, Vol. 83, No. 2, September 1977; Rogers, 2003.

[18] Markle Foundation, "The Markle Foundation Launches AI Initiative Focused on Emergency Preparedness, Response, and Recovery," December 10, 2025.

[19] Markle Foundation, "About Markle," webpage, undated.

[20] AIDE, homepage, undated.

[21] AIDE, undated.

[22] A *landscape assessment* is a systematic survey of a defined domain. It characterizes what exists, identifies patterns and gaps, and draws out implications from the analysis. It does not rank or evaluate products. Its purpose is to give stakeholders an independent picture of the terrain so that technology companies, funders, policymakers, and those supporting EM can have some understanding of where things stand in the field.

[23] Throughout the report, *today* and *current* mean at the time that data collection ceased in April 2026.

The Foundation's intention was to identify the extent to which products that address the most pressing EM needs exist and are positioned for broad adoption and diffusion across the field, particularly by SLTT governments and the nonprofit organizations that work alongside them. Whether those products are so positioned depends on both organizational adoption *and* field-level diffusion, as we described earlier in this chapter. With this remit, we set out to answer three research questions: (1) What AI-enabled products are currently on the market to support EM use cases? (2) What potential exists for the adoption and diffusion of these products to evolve EM practice? and (3) What must happen for this evolution to occur? In the rest of this report, we describe the AI-enabled product landscape and what it implies for achieving adoption and diffusion.[24]

The Scale and Novelty of the AI-Enabled Product Landscape Required a Structured, Mixed-Methods Approach

We conducted a four-stage mixed-methods study between October 2025 and March 2026. In stage one, we identified 45 EM task areas using authoritative sources (e.g., EM accreditation program standards, national planning frameworks) and validated those areas through expert review and engagement with emergency managers. In the second stage, we systematically searched for products aligned with these task areas using complementary strategies: (1) reviewing technology company lists from major EM and disaster conferences; (2) conducting two web searches, including iterative use of LLM application programming interfaces (APIs); (3) identifying relevant companies and products through commercial company data APIs; and (4) obtaining stakeholder and expert recommendations. This process yielded 1,892 candidate products, which we narrowed to 1,179 products from 717 companies. Stage three involved designing a framework to capture information that is decision-relevant for EM stakeholders and consistently available from public sources, drawing on established technology evaluation research. In stage four, we gathered and refined product information using LLMs and updated the framework as needed. We supplemented these efforts with engagement with technology companies, RAND quality assurance processes, and technical expert input to strengthen rigor. We then applied descriptive statistics and qualitative analysis to characterize the landscape and identify patterns. These results inform the analyses we present in Chapters 2 and 3. We provide additional detail in Appendix B.

What the Approach Could and Could Not Establish Was Shaped by the Nature of the Landscape

Our research approach carried meaningful limitations that readers should consider when interpreting the findings. Our assessment characterizes the AI-enabled products that surfaced from the search strategies we applied to the 45 task areas and not the universe of products that could exist in the EM space. Our findings about proportions, patterns, and concentrations describe the landscape we identified rather than support generalization beyond it. The analysis is based almost entirely on publicly available information, and characterizations requiring nonpublic information were not feasible to attain at scale. The use of LLM APIs as a strategy favored products with a strong online presence and clear AI-related language. Some product characterizations rest on inference from available evidence rather than direct technology company confirmation,

[24] We conducted this research independently. This report does not necessarily reflect the opinions of the Foundation or AIDE Initiative partner organizations. Attribution does not constitute and should not be interpreted as agreement, acceptance, or endorsement by the sponsor of the contents thereof.

particularly for connectivity requirements, deployment scale, and AI governance claims.[25] In this report, we characterize what products exist, how they are structured, and the conditions that may affect their adoption readiness; we could not identify what products are being used successfully in specific SLTT jurisdictions or nonprofits today or how those products perform. Information that was current at the time data collection concluded may already be outdated by the time this report is published. *Breadth came at the cost of depth on any given product across 45 task areas and 1,179 AI-enabled products from 717 technology companies, so readers who are seeking a deep comparative evaluation of specific products should treat this report as a starting point.* Cost was a dimension in which these limitations compounded: Pricing opacity, structural variation in access models, and the absence of standardized cost disclosure meant that the depth of cost analysis that was possible from public sources was limited by technology company practices and not by research effort. Appendix B details these limitations further.

We Share the Information That Led to This Conclusion Across Four Additional Chapters and Nine Appendixes

The remainder of this report presents the results of our efforts to answer this study's three research questions and the practical implications that emerge. In Chapter 2, we map the landscape of AI-enabled products with potential relevance to EM. In Chapter 3, we examine product features that could influence organizational adoption decisions and the extent to which the diffusion of AI-enabled products is achieved. In Chapter 4, we analyze the extent to which diffusion is likely in the near term across SLTT EM offices and other organizations that support EM functions. In Chapter 5, we translate those findings into a sequenced agenda for action by emergency managers and others who support EM functions, technology companies, policymakers, researchers, and others. We also introduce four high-impact ideas that stakeholders can implement in the next 12 months to propel progress toward the adoption and diffusion of AI-enabled products in EM, which concludes the report.

We provide supporting detail and reference material in the appendixes. In Appendix A, we provide a brief history of AI. In Appendix B, we describe the research design. In Appendixes C through E, we describe the response, recovery, and preparedness functional areas; the task areas within those functional areas; the number of products we identified to support each use case; and how AI enables the products at a high level, including the most common AI techniques that enable the products for each task area. In Appendix F, we identify gaps in the product market for certain task areas. In Appendix G, we provide sequenced action agendas for key stakeholder groups to support progress toward adoption and the diffusion of AI-enabled products. In Appendixes H and I, we provide tools that emergency managers and others who support EM functions can use as they get ready to implement AI-enabled products within their organizations.

[25] In the broadest sense, *AI governance* refers to the structures, rules, norms, and processes that shape how AI systems are developed, deployed, and used. AI governance can be reflected in laws, regulations, standards, organizational policies, and informal norms, as well as in other forms.

CHAPTER 2

Many Products Exist with the Potential to Support Emergency Management

In Chapter 1, we introduced the scope of the assessment, the 1,179 products identified, and the conditions under which AI-enabled products must be adopted to meaningfully strengthen EM. Now, we turn to the first of those topics: What products are in the landscape?

The overarching picture is that the landscape offers emergency managers and others who support EM functions real options for many of the task areas they work on, but the coverage that is offered is uneven overall, partial within most task areas, and marked by clear gaps. We describe what is in the landscape, where product coverage is most concentrated, and where it is partial or absent. Throughout this chapter, we offer examples of products to illustrate categories, clusters, or types of work.[1]

Many Products Are Purpose-Built for EM (or Have Been Used to Support It) and Are Hazard Agnostic

We examined the extent to which each product (1) was explicitly designed and marketed for one or more task areas (hereafter termed *purpose-built*); (2) was demonstrated through evidence to have been used in an EM context with core AI features that are highly relevant to one or more functional areas and task areas; or (3) was a general AI-enabled platform configurable for an EM functional area and one or more task areas—not purpose-built but applicable.

Approximately half of the products we identified were either purpose-built for EM or had been used to support EM in a specific case study or deployment (e.g., use in a named disaster event or a customer report describing such use) (Figure 2.1). Some technology companies have products that span multiple levels: For example, Salesforce offers an emergency response management solution, designed specifically for EM, while also marketing Agentforce for Energy and Utilities for storm response.[2]

For each of the 45 task areas, the percentage of products that are purpose-built for EM varies substantially.[3] Fire Management and Suppression is proportionally the most purpose-built, and 73 percent of products in

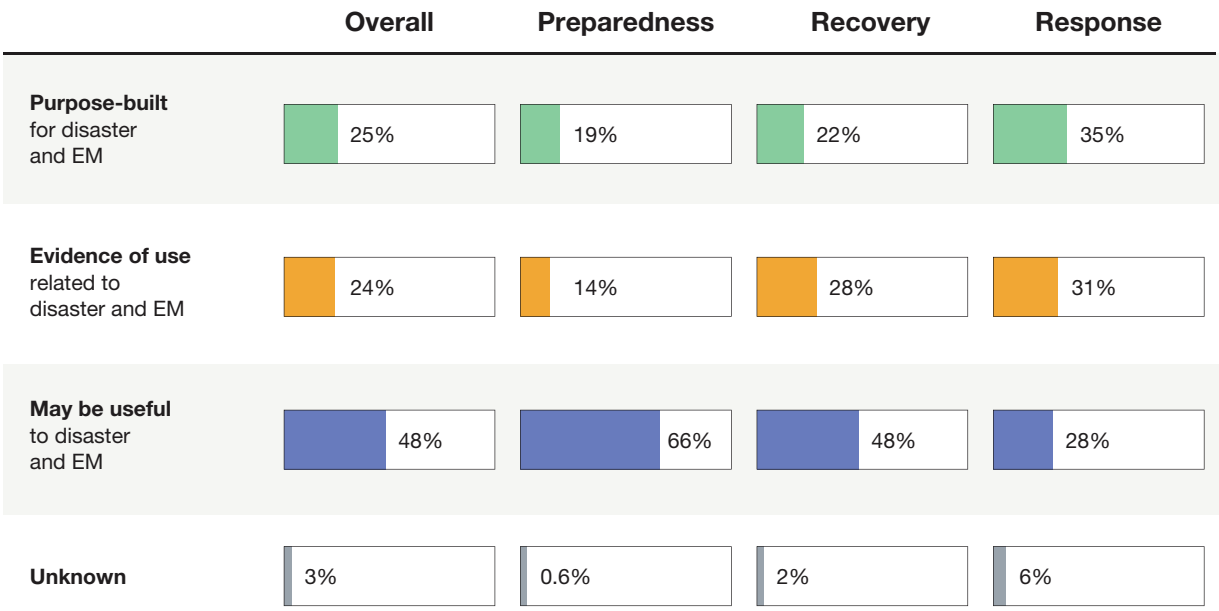
[1] These examples are reference points drawn from the landscape we assessed, selected because our analysis suggests that they represent the category being described. They are not evaluations, comparisons, or endorsements of product performance. Any number of other products in the landscape could have served as examples in the same places, and a product used to illustrate one category could in many cases have been used to illustrate others. The AI-enabled product market moves rapidly, and our characterizations draw on public information, much of which came from technology companies. Individual product characterizations could be imperfect as a result. Despite this, we judged that specific examples were useful to the reader as reference points and have included them throughout with that purpose in mind.

[2] Salesforce, "Emergency Response Management," webpage, undated-b; Salesforce, "Agentforce for Energy & Utilities," webpage, undated-a.

[3] See Appendixes C through E for the full list of task areas.

FIGURE 2.1

The Extent to Which Products Are Purpose-Built for EM Overall and by Functional Area



NOTE: Percentages may not sum to 100 percent because of rounding. *Unknown* reflects products for which EM alignment could not be determined.

that task area are designed specifically for EM use. Situational Awareness, Hazard Monitoring, and Information Sharing has the highest number of purpose-built products, at 241 products (or roughly 45 percent). The percentage of purpose-built products for most task areas sits well below those two. Twenty-one task areas had ten or fewer purpose-built products, and several task areas had a very low proportion of such products. One task area, Natural and Cultural Resources Restoration, had zero purpose-built products.[4]

No product category is inherently better: mission alignment is not the same as technical capability, and a purpose-built product is not automatically a better fit than a configured general platform. This distinction does not, in and of itself, indicate how much value a product may offer an EM customer. Rather, it provides a starting question for anyone who is exploring these products: Is this product built around EM, adapted to EM, or being offered to EM as a configurable general capability?

More than 80 percent of the products are *hazard-agnostic*, meaning that they do not require per-hazard configuration, training, or module development. Another 5 percent are products that were purpose-built for one hazard, but the underlying technology is inherently capable of addressing adjacent hazards, and the technology companies have already begun or have explicitly signaled such an expansion. We refer to these products as *single-adaptable* (e.g., ResiQuant, Tenax ai, ALERTCalifornia, Google SOS Alerts, Urbint Storm Impact). Of the 15 percent of hazard-specific products that are available, wildfire is the hazard that is most deeply served.

[4] Some products that support Natural and Cultural Resources Restoration began as general environmental monitoring products and have since been applied in disaster contexts.

There Are Nine or More Products Available to Support Every Functional Area and Task Area

The market heavily favors response (53 percent of products), serves recovery (36 percent), and thinly covers preparedness (12 percent).[5] Our analysis shows that most products are associated with only one functional area, but about 21 percent are associated with two functional areas (typically response and recovery). Only 0.6 percent supported all three functional areas.

All functional areas and all task areas within them have nine or more associated products. Within response, the market is concentrated in a few dominant task areas. More than 75 percent of response products support either Situational Awareness, Hazard Monitoring, and Information Sharing or Incident Coordination and Operational Decision Support. Approximately 34 percent of response products support both of these task areas. The products associated with recovery are far more distributed. Damage and Impact Assessment, the recovery area with the most products, accounts for 22 percent of products in this area. Product support across preparedness task areas is thin but evenly spread out: No single preparedness task area is well served, but no task area is completely empty either (Figure 2.2).

For EM customers, this means that the quantity of potential products varies across functional areas. They will have many response-supporting products to choose from, fewer choices for recovery, and far fewer choices for preparedness. For technology companies and others supporting EM, this disparity in product availability highlights opportunities for growth and innovation.

These results also suggest that one cannot assume that a product evaluated for response will be useful for recovery and that a product evaluated for preparedness will be useful for response. The extent to which that matters depends on the purchasing entity's resources and preferences (e.g., whether the entity prefers integrated products across functional areas). Customers will need to look past marketing labels and technology company descriptions.

Most Products Support Multiple Task Areas

Individual products support from one to 13 task areas. The products vary significantly in the number of task areas they support (Figure 2.3), but 78 percent support two or more task areas and 53 percent support three or more.[6] This finding may result from intentional design choices by technology companies or partly from inherent overlap between the task areas that blur the boundaries between them.

The two task areas that most often co-occur in products both fall under response: Situational Awareness, Hazard Monitoring, and Information Sharing and Incident Coordination and Operational Decision Support. The next most frequent pairings combine Initial Damage and Impact Assessment in response with Data, Information Management, and Situational Awareness in recovery. We identified only 15 products that support nine or more task areas. All products in this subset support more than one functional area, and nearly all span recovery and response. Eighty percent of the products in this subset were purpose-built for EM.

[5] Percentages may not sum to 100 percent because of rounding.

[6] This finding may in part reflect the way that the task areas overlap in operational practice.

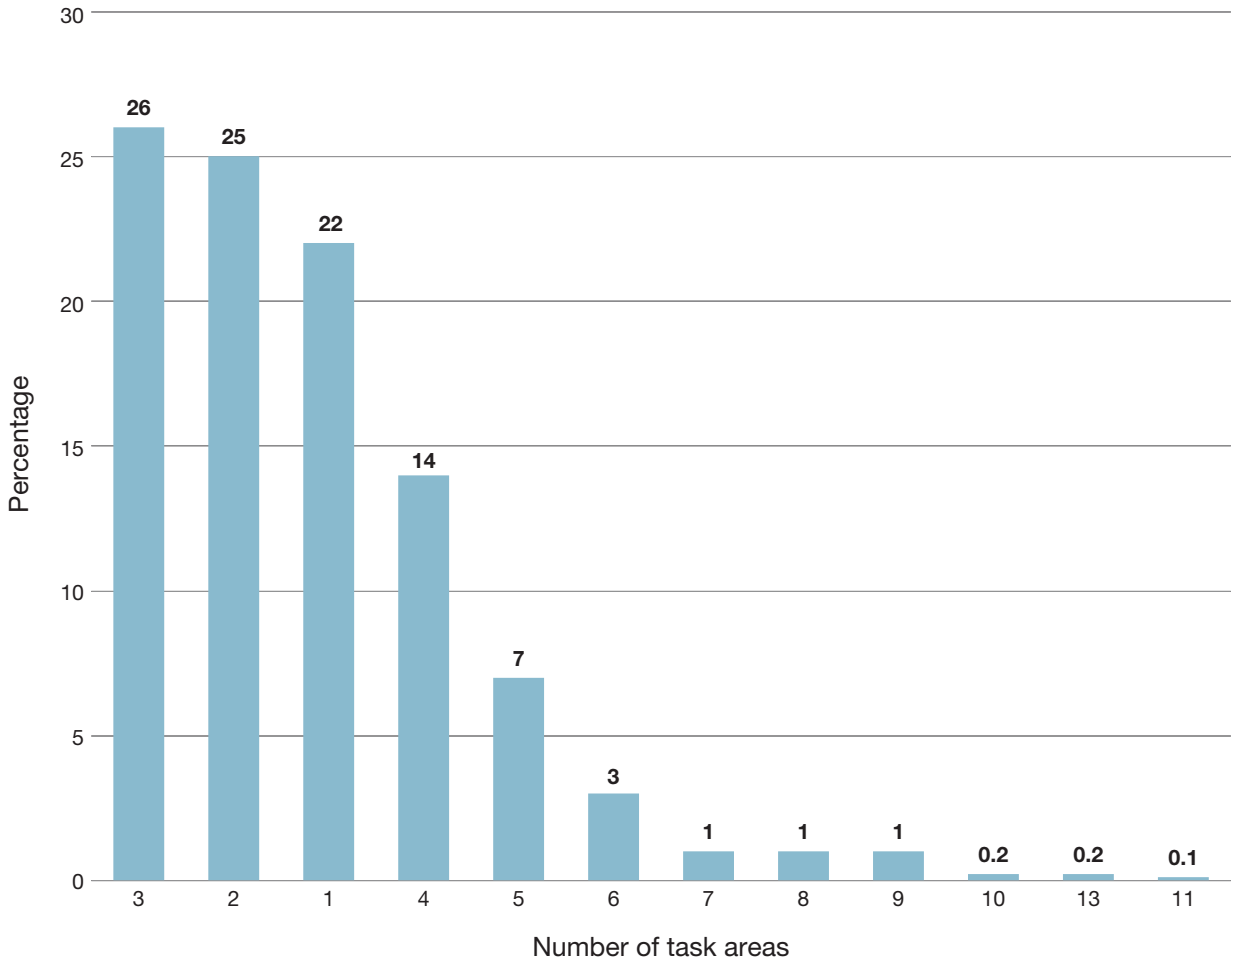
FIGURE 2.2

Task Areas with the Highest and Lowest Number of Products by Functional Area

	Response	Recovery	Preparedness
↑ **Top five**	Situational Awareness, Hazard Monitoring, and Information Sharing	Damage and Impact Assessment	Risk Assessment
	Incident Coordination and Operational Decision Support	Data, Information Management, and Situational Awareness for Recovery	Capacity and Capability Analysis
	Initial Damage and Impact Assessment	Financial Recovery, Grants, Cost Tracking, and Insurance	Hazard Identification and Vulnerability Assessment
	Public Information, Alerts, and Protective Guidance	Community Engagement, Communication, and Information Access	Resource Management, Mutual Aid, Logistics, Distribution, and Supply Chain Resilience
	Resource Request, Allocation, and Logistics	Infrastructure and Lifelines Restoration	Training
Bottom five ↓	Fatality Management Services	Natural and Cultural Resources Restoration	Program Governance, Authorities, and Administration
	Emergency Debris Clearance and Route Opening	Private-Sector and Public-Private Partnership Recovery Coordination	Exercising and Continuous Improvement
	Family Reunification and Welfare Checks	Volunteer, Donations, and Community Support Coordination	Integrated Preparedness Planning and Plan Synchronization
	Volunteer, Donations, and Offers Management	Debris, Waste, and Environmental Safety Management	Community Engagement and Whole Community Partnerships
	Shelter, Mass care, and Human Services	Land Use, Rebuilding, and Mitigation Planning Through Recovery	Training

NOTE: Training appears twice because there are only nine preparedness task areas.

FIGURE 2.3
Number of Task Areas That the Products Support



NOTE: Products supporting more than 13 task areas are not shown (<0.1 percent). Percentages may not sum to 100 percent because of rounding. Only one product supported 12 task areas, and it is not depicted.

Products Supporting Task Areas That Are Important to Emergency Managers Exist

The task areas most heavily covered by AI-enabled products tend to be those in which EM offices are both the primary user and the primary purchaser, which gives technology companies a clearer market to build for. In the following sections, we describe four of those task areas, including how AI tends to enable products within those areas. In Appendixes C through E, we characterize at a high level how AI enables products in each task area and what AI techniques we most frequently observed in products within those task areas.

Products That Support Situational Awareness, Hazard Monitoring, and Information Sharing

Situational Awareness, Hazard Monitoring, and Information Sharing is the single most densely populated task area in the landscape, and the products serving it use the widest variety of AI techniques. At a high level, AI in these products collects and analyzes data about hazards, events, and conditions at scale. Computer

vision (CV) and multimodal AI track changes in the environment from satellites, drones, cameras, and sensors. Machine learning (ML) generates forecasts and detects anomalies in large datasets. NLP extracts signals from social media, news, and other text sources. The products vary from raw data providers, to analytic layers that classify what the raw data shows, to decision-support dashboards that aggregate multiple feeds for operators.

A pattern cuts across the task area: Situational Awareness, Hazard Monitoring, and Information Sharing functions as a prerequisite input to nearly every other task area, which means that an EM office's choices in this task area shape what downstream products it can use effectively. The *sharing* component of the task area is often not a separate product but a product feature delivered through APIs, dashboards, and geographic information system (GIS) integrations that pass data to other products. Six clusters of products carry the most operational weight for EM offices. Table 2.1 summarizes what the products in each cluster do and lists example products.

Additional clusters serve specific hazards or specific sensing domains and may be central for EM offices depending on the context, such as wildfire camera networks, Internet of Things (IoT) and environmental sensors, flood and water monitoring, air quality, and maritime awareness.

Products That Support Incident Coordination and Operational Decision Support

Incident Coordination and Operational Decision support is the second-most densely populated task area in the landscape after Situational Awareness, Hazard Monitoring, and Information Sharing. However, the high density of this task area partly reflects products with a coordination function that is secondary to their primary capability, along with general-purpose AI platforms. At a high level, AI in these products structures

TABLE 2.1

Situational Awareness, Hazard Monitoring, and Information Sharing: Product Clusters, What Products Do, and Example Products

Product Cluster	What Products Do	Example Products
Satellite and aerial imagery	Provides pre- and post-event imagery, AI-derived change detection, and persistent monitoring	• Planet Lab's Planet Monitoring • BlackSky Spectra • The Vantor Hub
Weather intelligence	Applies ML to model and observational data to generate hyperlocal forecasts, impact scores, and sector-specific alerts	• DTN WeatherSentry • Tomorrow.io • WeatherOptics
Drone and uncrewed aerial system situational awareness	Provides live aerial video, thermal imaging, AI object detection, and mapping from crewed or autonomous aircraft	• Skydio X10 for Drone as First Responder • DJI Matrice 4T • BRINC Responder
Operations centers and data fusion	Integrate CAD, video, sensor, and other feeds into a unified command view	• Axon Fusus • Peregrine Technologies • Palantir Foundry and Gotham
OSINT and threat monitoring	Scans social media, news, and open sources at scale using NLP to detect emerging events and narrative risks	• Dataminr Pulse • Babel Street • Zignal Labs
Critical event management	Aggregates multisource risk data into dashboards, correlates threats to assets and people, and issues targeted alerts	• Everbridge 360 • OnSolve • AlertMedia

NOTE: CAD = computer-aided dispatch; OSINT = open-source intelligence.

the flow of operational information and decisionmaking during an incident. ML generates predictions and correlations across operational data. NLP and ML applied to generative AI supports conversational queries and summarization and drafts briefings and situation reports. Knowledge representation encodes incident coordination workflows and protocol logic. Planning, scheduling, and optimization techniques support task sequencing and resource assignment.

A pattern cuts across this task area: Products in this area span platforms that purport to structure the entire incident management life cycle to point solutions that provide specific decision inputs. AI decision assistants are also starting to emerge: For instance, we observed some AI decision assistants embedded in emergency operations center (EOC) platforms. Within these platforms, the AI assistants summarize board data, generate situation reports, and answer queries in natural language. This emergence represents AI moving from providing data to providing synthesized judgment support. Four product clusters carry the most operational weight for EM offices. Table 2.2 summarizes what the products in each cluster do and lists example products.

Weather decision support is also part of this task area; the cluster of weather intelligence products that provide both situational awareness and operational decision inputs is covered in the Situational Awareness, Hazard Monitoring, and Information Sharing section above.

Additional product clusters serve specific hazards, responder domains, or coordination functions and may be central for EM offices depending on the context, such as computer-aided dispatch systems, fire and hazard decision support, drone fleet management, satellite tasking, and AI-based scheduling and planning.

Products That Support Damage Assessment

Damage and impact assessment appears twice in the task area framework used for this assessment: once during response (Initial Damage and Impact Assessment) and once during recovery (Damage and Impact Assessment). The two task areas share roughly 40 percent of their products, and the distinction between these task areas is temporal and operational rather than technological. For that reason, we treat them together here. At a high level, AI in these products converts imagery, whether from satellites, aircraft, drones, vehicles, or ground sensors, into structured damage data, such as FEMA damage categories, severity ratings, and counts of affected structures. CV does the central work of extracting damage features from the imagery. ML classi-

TABLE 2.2
Incident Coordination and Operational Decision Support: Product Clusters, What Products Do, and Example Products

Product Cluster	What Products Do	Example Products
EM platforms	Provide configurable workflows, status boards, dashboards, and ICS products for EOC operations	• WebEOC Nexus • COBRA 5 • eICS
Critical event management	Coordinates multisite and multiagency responses through risk correlation, automated workflows, and communications	• Everbridge 360 • OnSolve • AlertMedia
AI decision assistants	Answer natural-language questions against operational data, generate briefings, and summarize incidents	• Juvare AI Assistant • PRATUS • COBRA KAI
Resource visualization	Displays real-time positions of personnel, vehicles, and assets on maps	• FLORIAN • FleetUp • MARVLIS

NOTE: ICS = incident command system.

fies damage levels and generates risk scores. Multimodal AI fuses optical imagery, synthetic aperture radar (SAR), and sensor data. Generative AI drafts reports from field data.

A pattern cuts across the task area: Products in this area span the full sensing range, from satellites, which offer large-area coverage at lower resolution, to drones, which offer site-specific, high-detail imagery, to ground sensors, which offer building-specific continuous monitoring. An asymmetry sits beneath this range. Products that are capable of operating during active response inherently produce artifacts that persist into recovery, but products built for recovery sometimes cannot operate under response conditions. Recovery workflows often require conditions that are unavailable during an active event, such as safe interior property access for detailed surveys or post-event data that does not yet exist during response. *The practical implication for EM offices is that a response-capable damage assessment product generally provides recovery value too, while the reverse is not true.* Five product clusters carry the most operational weight for EM offices. Table 2.3 summarizes what the products in each cluster do and lists example products.

Additional product clusters serve specific infrastructure types, specific hazards, or specific downstream users and may be central for EM offices depending on the context, such as road and infrastructure damage detection, structural health sensors, flood damage estimation, light detection and ranging mapping, and insurance and parametric assessment.

TABLE 2.3

Damage Assessment: Product Clusters, What Products Do, and Example Products

Product Cluster	What Products Do	Example Products
Satellite damage detection	Compares pre- and post-event satellite imagery using AI to classify building damage severity at scale	• CrowdAI • DamageSense • xView2
Aerial and drone damage survey	Captures high-resolution imagery from drones and aircraft to generate damage maps, 3D models, and orthomosaics	• EagleView Disaster Response Program • DroneDeploy Aerial • Pointivo
Property-level risk and damage scoring	Applies CV to aerial imagery for property-specific damage ratings, roof condition, and risk scores	• CAPE real estate property intelligence • Betterview • ZestyAI
Field data collection	Enables inspectors to capture geo-referenced damage data on mobile devices with AI-assisted form filling	• ArcGIS Survey123 • Fulcrum • Felt field data collection app
AI damage classification models	Employ pretrained deep learning models that classify FEMA damage levels from imagery	• ArcGIS Damage Assessment (Drone Imagery) model • xView2 • Humanitarian OpenStreetMap fAIr

Products That Support Public Information, Alerts, and Protective Guidance; Protective Action Decision and Implementation; and Evacuation and Transportation Operations

The Public Information, Alerts, and Protective Guidance; Protective Action Decision and Implementation; and Evacuation and Transportation Operations task areas are operationally interrelated. A protective action decision triggers an alert, an alert directs people to take an action, and an evacuation operation moves them. Products in these three task areas often work in sequence, and some technology companies build across more than one of them. At the same time, the AI work that each task area requires is distinct. Alerts are dominated by NLP and generative AI for message composition and translation, and ML in the threat detection triggers alerts. Protective action decisions are dominated by predictive modeling of hazard progression, CV for hazard detection and planning, and optimization techniques for evacuation sequencing. Evacuation operations are dominated by predictive modeling of traffic and clearance times and planning and optimization for routing and signal control. For these reasons, we are treating the three task areas together in this section but separating the product market for each task in the tables below.

A pattern cuts across the three task areas: AI is concentrated at the front end (i.e., trigger detection, hazard modeling, content generation, route optimization) and at the analytic end (i.e., traffic modeling, mobility analysis), while the execution layer (i.e., alert delivery, order issuance, physically moving people) remains largely rules-based or human-driven. *An EM office evaluating products in these task areas will generally find that AI provides inputs to decisions or communication rather than automates the operations.*

Public Information, Alerts, and Protective Guidance

At a high level, AI in these products detects conditions that warrant alerting, generating, and translating alert content and surfaces authoritative information to the public. Six product clusters carry the most operational weight for EM offices. Table 2.4 summarizes what the products in each cluster do and lists example products.

TABLE 2.4

Public Information, Alerts, and Protective Guidance: Product Clusters, What Products Do, and Example Products

Product Cluster	What Products Do	Example Products
Mass notification platforms	Deliver geo-targeted alerts across short message service, voice, email, mobile push, sirens, signage, and social media, often integrated with the integrated public alert and warning system and wireless emergency alerts	• Everbridge Public Warning • Genasys ALERT • CodeRED
AI message drafting	Uses generative AI to compose, optimize, and translate emergency notification text rapidly	• AlertMedia AI Assistant • COBRA KAI • AlertAI
Translation and multilingual delivery	Provides real-time translation of alert content into dozens of languages	• Wordly • LILT • Google Cloud Translation
Public-facing hazard information	Surfaces authoritative emergency data directly to consumers via search, maps, or apps	• Google SOS Alerts • Watch Duty • MyRadar
Emergency broadcast	Auto-captures government alerts and rebroadcasts via radio, TV, and streaming with AI text-to-speech in multiple languages	• BEACON
Facility-level alerting	Integrates sensors, such as gun detection and environmental monitors, with in-building notification devices	• XSponse • Quicklert QBOX and ConnecTab • ZeroEyes

Within this task area, most mass notification platforms use ML primarily in their risk intelligence layer, while the message delivery is rules-based. Generative AI is the newest AI technique entering this space, specifically supporting rapid message composition.

Protective Action Decision and Implementation

At a high level, AI in these products models how a hazard will progress so human decisionmakers can determine the timing and scope of protective actions. Three product clusters carry the most operational weight for EM offices. Table 2.5 summarizes what the products in each cluster do and lists example products.

The actual implementation of protective action (e.g., issuing orders, moving people) typically relies on products that are mapped to the Public Information, Alerts, and Protective Guidance and Evacuation and Transportation Operations task areas. AI is concentrated in the decision input rather than in the act of implementing the decision.

Evacuation and Transportation Operations

At a high level, AI in these products predicts traffic behavior, optimizes routes and signals, and analyzes mobility patterns to support the movement of people during an evacuation. Five product clusters carry the most operational weight for EM offices. Table 2.6 summarizes what the products in each cluster do and lists example products.

Additional product clusters across the three task areas serve specific coordination functions or specific responder groups and may be central for emergency managers depending on the context (e.g., traffic clearance analysis and lockdown and shelter-in-place automation in protective action decisions).

Some of the products described in the previous sections are new, but there are many AI features that technology companies are rapidly adding to existing products. Established products that are central to GIS, broad EM platforms, and mass notification all sit in this category. For emergency managers, the term *AI products* often means that AI capabilities have been layered onto purpose-built products that they are already aware of, have used, or currently use rather than having to contend with entirely new products. Additionally, emergency managers may already have relationships with the associated technology companies.

TABLE 2.5

Protective Action Decision and Implementation: Product Clusters, What Products Do, and Example Products

Product Cluster	What Products Do	Example Products
Evacuation zone management	Predefines and activates zone-based evacuation orders with public-facing maps and multichannel notification	• Genasys EVAC and Protect • Perimeter Map
Fire spread modeling for decisions	Simulate fire growth scenarios to determine when and where to issue evacuation orders or shelter-in-place instructions	• Technosylva Wildfire Analyst (WFA-e) and FireRisk • EmberCast • Ladris Fire
Flood forecast-to-action	Translates flood depth and timing forecasts into protective action recommendations	• Google Flood Hub • aquaFLOODFirst Street Flood Factor

TABLE 2.6
Evacuation and Transportation Operations: Product Clusters, What Products Do, and Example Products

Product Cluster	What Products Do	Example Products
Evacuation traffic modeling	Predicts congestion, clearance times, and choke points for planned or ad hoc evacuation zones	• Genasys Protect TRAFFIC AI powered by Ladris Perimeter Plan • AnyLogic
Route optimization	Calculates optimal routes for evacuation vehicles and responders considering real-time conditions	• INRIX AI Traffic • Samsara routing and dispatch • WeatherOptics RightRoute
Transit operations	Help transit agencies add or cancel trips, reroute buses, and push real-time passenger information	• Swiftly Service Adjustments and real-time passenger information • Via Intelligence
Mobility analytics	Analyze movement patterns to identify evacuation bottlenecks and populations with limited egress	• StreetLight InSight • Spectus • Cuebiq
Mapping and navigation	Provide real-time route guidance, including road closures and disaster overlays	• Google Maps • Waze

There Are Products Other Entities That Support EM Functions Would Use

The task areas that anchored this research were not limited to task areas performed solely by personnel in SLTT EM offices. Although some task areas represent work done primarily by such personnel, others involve the broader distributed function, and many involve both. The product landscape assessed in this report therefore reflects products that might support EM work even though the user of the product might be any entity that supports EM. We identified hundreds of products of this type. Given the breadth of potential end users and the wealth of available products they might use, we highlight in this section subsets of products that might be most used by utilities, nonprofits, or the public. The sections that follow highlight how AI enables the products that these three stakeholder groups might use to support their contributions to EM functions. Appendixes C through E characterize at a high level how AI enables products in each task area and what AI techniques we most frequently observed in products within the task area.

Products That Are Oriented to the Task Areas That Utilities Support

Utilities, including electric, water and wastewater, gas, and telecommunications providers, are responsible for maintaining or rapidly restoring the essential services that communities depend on during and after disasters. Critical Infrastructure and Essential Services Stabilization is the task area under response in the framework used for the assessment that captures this work, and utilities are likely to be the primary users of products that fall under this task area.

At a high level, AI in these products predicts infrastructure behavior under stress and optimizes the operation of infrastructure systems before, during, and immediately after disaster events. ML predicts outages, assesses asset health, and forecasts demand. Planning, scheduling, and optimization techniques support the dispatch of distributed energy resources and the balancing of grid loads. CV supports the rapid inspection of sewer and pipeline conditions. Knowledge representation encodes grid control rules and fault diagnosis logic.

A pattern cuts across this task area: AI applications in utilities cluster around three kinds of work. The first predicts what infrastructure will do next, including outages, equipment failures, and demand surges. The second controls distributed and complex systems, including grid operations, microgrid dispatch, and virtual power plant aggregation. The third rapidly assesses damage and identifies workarounds after an event, par-

ticularly for water, sewer, and pipeline infrastructure. Five product clusters carry the most operational weight for utilities. Table 2.7 summarizes what the products in each cluster do and lists example products.

Additional product clusters serve specific infrastructure types or specific resilience functions and may be central for utilities depending on the context, such as virtual power plant aggregation, sewer infrastructure assessment, cooling for critical facilities, pipeline monitoring, and underground utility monitoring.

TABLE 2.7

Critical Infrastructure and Essential Services Stabilization: Product Clusters, What Products Do, and Example Products

Product Cluster	What Products Do	Example Products
Grid management and ADMS	Orchestrate distribution grid operations, including fault detection, outage management, and distributed energy resource dispatch	• GE Vernova GridOS ADMS • AutoGrid Flex Saver • Camus Energy
Outage prediction	Forecasts storm-driven power outages days ahead to support pre-staging of crews and equipment	• DTN Storm Impact Analytics • IBM Outage Prediction • AiDash Climate Risk Intelligence System
Microgrid control	Manages distributed energy resources (e.g., solar, batteries, generators) to provide resilient local power when the main grid fails	• Generac ARC • PXiSE • GridPoint
Water and wastewater stabilization	Monitors and controls water treatment and distribution in real time, including contamination detection	• Xylem Vue • KETOS smart water intelligence
Operational technology cybersecurity	Protects critical infrastructure control systems from cyberattack	• Claroty xDome and Continuous Threat Detection • CrowdStrike Falcon

NOTE: ADMS = advanced distribution management system.

Products That Are Oriented to the Task Areas Nonprofits Support

Many communities, nonprofits, members of the National Voluntary Organizations Active in Disaster coalition, long-term recovery groups, and faith-based organizations do much of the direct work of helping with mass care during the response to and long-term recovery of individuals and households. Nonprofits are likely the primary users of products in both the Volunteer, Donations, and Offers Management task in response and the Volunteer, Donations, and Community Support Coordination task in recovery in the framework used for this assessment. This subsection focuses on both tasks.

Volunteer, Donations, and Offers Management and Volunteer, Donations, and Community Support Coordination

AI in these products primarily supports the skills-based matching of volunteers to opportunities, tasks, and activations during response using ML; planning and optimization for scheduling; NLP for communications; and generative AI for drafting content. Overall, the greatest value lies in matching rather than screening or vetting (Table 2.8). In recovery, the support shifts to sustained coordination over months—matching volunteers to specialized tasks (e.g., case management, construction). Effectiveness is dependent on preestablished, credentialed volunteer rosters and supported by such features as credential verification and exportable records for FEMA reimbursement and grant matching.

TABLE 2.8

Volunteer and Donations Management: Product Clusters, What Products Do, and Example Products

Cluster	What Products Do	Example Products
General volunteer management	Recruits, screens, schedules, and tracks volunteers with AI-optimized matching	• Galaxy Digital Get Connected • Bloomerang Volunteer • Volgistics
Volunteer skills matching	Aligns volunteer skills and availability with specialized recovery tasks	• Taskade AI volunteer matching agent • Golden intelligent matching system
Humanitarian mapping	Coordinates volunteer mappers to create baseline geographic data for disaster response using human-AI collaboration	• Humanitarian OpenStreetMap Tasking Manager • Humanitarian OpenStreetMap fAIr
Corporate volunteering	Matches employees to volunteer opportunities aligned with skills and interests	• Benevity • StratusLIVE Community Volunteer Center
Fundraising and donor CRM	Predict donor propensity, automate outreach, manage donor relationships, and convert volunteers to donors	• Bonterra EveryAction • DonorSearch Ai • Virtuous CRM+
Online giving platforms	Handle transactional fundraising, including forms, payment processing, ask amount optimization, and fraud detection	• Fundraise Up • Pledge.ai • DonationXchange and CharityAI
In-kind donation matching	Matches surplus corporate goods and other material donations to nonprofit needs	• Good360 • Agentforce • NeedsList and RespondLocal
AI-powered donor advising	Recommends vetted disaster relief charities to donors	• BBB's AskGive

NOTE: CRM = customer relationship management.

Relatedly, AI-enabled donations management products match financial and in-kind contributions to needs, campaigns, and organizations, combining transaction-layer optimization (e.g., ask amounts, fraud detection) with relationship management (e.g., donor segmentation, engagement, and retention) (Table 2.8). As with volunteer management, effectiveness depends on sustained engagement over time, and in-kind matching and donor coordination becomes particularly important in the recovery phase.

Products That Are Oriented to the Public

Dozens of products in the landscape are designed primarily for the public rather than for emergency managers or the organizations that support EM work. These products can improve household preparedness, support safer response decisions, and speed recovery, and they also touch on task areas, including Public Information, Alerts, and Protective Guidance; Individual and Household Assistance and Case Management; Evacuation and Transportation Operations; Community Engagement, Communication, and Information Access; and Community Engagement and Whole Community Partnerships. Many of these products are not yet universally known or adopted across the United States. The products span a variety of functions and include

- products that support individuals and households doing home inventories (e.g., Bevel)
- AI chatbots that help survivors and the public locate disaster assistance (e.g., SOS Connect, Just Ask Mia, the Red Cross' Clara)
- multi-hazard warning and situational awareness applications (e.g., Disaster Alert, Watch Duty)
- crisis information surfacing products (e.g., Google SOS Alerts).

These products reach the public through diverse channels and access models, from free downloadable apps, to embedded features in mainstream products, to partnership-activated ride programs. The AI role in these products varies considerably, both in what AI does and how specifically the AI role is described in available product materials.

Nearly Half of the Products Identified in This Research Are General Solutions That Have the Potential to Support EM Task Areas

Nearly half of the products that we identified in this research are not purpose-built for EM. Rather, they are horizontal platforms and general-purpose products that serve many industries and domains with varying functional area applicability or disaster-specific features, workflows, or terminology. These products were included because we cast a broad net that captured products that had recognizable AI capabilities and could be logically connected to at least one defined EM task area retained in the taxonomy.

The defining characteristic of this category of products is theoretical applicability without evidence of a technology company's intent to use the technology in EM applications. This means that the product has identifiable AI capabilities and appears to be applicable to a task area but neither the technology company nor any credible source we could find connected the product formally to EM in a substantive way. Examples include customer relationship management platforms, contact center products, translation services, analytics dashboards, grant-writing assistants, workforce development platforms, and supply chain management suites. See the box on the next page for a discussion of one of these examples.

The result of casting such a broad net is a large proportion of wide and internally heterogeneous products within the total we identified. Products in this space vary, from enterprise software that exists entirely for other purposes but could conceivably be helpful if configured appropriately, to products with more obvious but still indirect relevance. The concentration of general-purpose products is particularly pronounced in

Translation Services: A Type of General Product That Is Applicable to EM Task Areas

Translation services products apply AI, typically neural machine translation combined with NLP, to convert written text and spoken language and document content from one language to another. Capabilities across the category include real-time translation of live conversations, meetings, and broadcasts; document translation that preserves formatting and terminology; translation integrated into chat platforms, email, and collaboration tools; and translation paired with captioning for accessibility. Several products offer optional human review or human-in-the-loop editing for higher-stakes communications. These products could reduce the time and effort required to translate public information and alerts, meeting content, inspector notes, and survivor communications. Example products include Google Cloud Translation, Microsoft Translator, Amazon Translate, LILT, Ai Translate by LSI, Boostlingo, and KUDO.

preparedness. Within this broader group, several subcategories of products are worth distinguishing by the type of work they support (Table 2.9).

General-purpose LLMs could occupy a distinct and cross-cutting position in the EM technology landscape. These products are not purpose-built EM products but rather platforms that those doing the work of EM could use to potentially accelerate their efforts. Examples of such commercial products caught in our net include ChatGPT (OpenAI), Gemini (Google), Claude (Anthropic), and Meta's Llama.

LLMs are trained in a two-stage process: First, they are trained on broad internet-scale text to develop general language understanding, and then they are fine-tuned through reinforcement learning from human feedback and other methods to improve helpfulness, accuracy, and safety. The result is a system that can engage in multi-turn conversation, retain context within a session, follow complex instructions, and produce outputs in virtually any written format. Multimodal versions (e.g., OpenAI's GPT, Google's Gemini) also process images, audio, and documents as input. Vendors continuously update and improve models, meaning that capability and accuracy evolve over time.

LLMs are widely accessible through browser-based and enterprise offerings, and major providers offer free, low-cost, or subscription access through conversational web interfaces that require little or no technical integration for baseline use.[7] A principal source of their value is broad general-purpose capability across many language and reasoning tasks rather than optimization for a single domain-specific workflow.[8] Base LLM deployments do not, by default, have authenticated access to authoritative real-time data sources, verified agency databases, or operational systems unless they are specifically connected through retrieval tools, APIs, or enterprise integrations.[9] Their outputs are probabilistic and can include inaccurate or fabricated

[7] Nestor Maslej, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njenga Kariuki, Emily Capstick, Anka Reuel, Erik Brynjolfsson, John Etchemendy, et al., *Artificial Intelligence Index Report 2025*, AI Index Steering Committee, Institute for Human-Centered Artificial Intelligence, Stanford University, April 2025, Chapter 4.

[8] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al., "Language Models Are Few-Shot Learners," in Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, eds., *Advances in Neural Information Processing Systems*, Vol. 33, 2020; Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al., "On the Opportunities and Risks of Foundation Models," arXiv, arXiv:2108.07258, 2021.

[9] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-Tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, *Advances in Neural Information Processing Systems*, Vol. 33, 2020.

TABLE 2.9
General Tasks That AI-Enabled Products Might Support, Product Type Descriptions, and Example Products

General Task	Product Description	Example Products
Meetings and communication	Products in this area apply generative AI and NLP to the coordination and communication layer of EM work. These products automatically generate meeting summaries and action items, draft emails and situation reports, surface relevant information from prior conversations and documents, and could help staff move faster through the routine written communication that accompanies activations, planning cycles, and stakeholder coordination. During a major activation, when staff are stretched thin, the time savings from automated notetaking, drafting, and document synthesis can be material.	• Zoom AI Companion or Microsoft 365 Copilot document processing • Workflow automation AI
Document processing and workflow	Document processing and workflow products address the data-intensive, form-heavy side of EM operations. AI is used to extract structured data from forms, pdfs, handwriting, and images (e.g., damage assessments, reimbursement documentation, survivor intake) and to enable natural-language search and Q&A queries across internal documents, such as plans, SOPs, and AARs. These systems also support case management, resource tracking, and assistance workflows while enabling users to build apps, automate processes, and generate dashboards through low-code, natural-language interfaces, reducing reliance onIT.	• Amazon Textract • Azure Document Intelligence • Azure AI Search with Azure OpenAI Service • Microsoft Dynamics 365 • Microsoft Power Platform with Copilot
Enterprise AI assistants (sometimes called copilot-style products)	Enterprise AI assistants (often called *copilot* products) add a conversational layer to existing platforms, letting users interact with data and functions through natural language instead of menus or code. The AI interprets user prompts, retrieves and synthesizes information from connected data sources, generates outputs (e.g., summaries, reports, drafts), and automates workflows by translating plain language into structured actions. In business intelligence products, such as Microsoft Power BI, this functionality enables the creation of visualizations and analyses from prompts; in customer service platforms, such as Zendesk, it supports ticket triage, response drafting, and case summarization. Across use cases, these systems operate on organization-provided data, so privacy risk depends on the sensitivity of connected data sources.	• Amazon Q Business • Microsoft Copilot in Power BI • Zendesk Copilot

NOTE: AAR = after-action report; IT = information technology; SOP = standard operating procedure; Q&A = question and answer.

content (e.g., hallucinations or confabulations), making validation and human review especially important for consequential uses.[10] Furthermore, researchers have not yet explored the variety of use cases to which these LLMs are appropriately and effectively applied, and guidance stemming from that research is not yet available to support emergency managers.

Beyond conversational interfaces, the same underlying models increasingly power *agentic AI products*, which are systems that can operate with greater autonomy to plan, call tools, retrieve information, write or modify code, automate workflows, and execute multistep tasks.[11] Examples include Claude Code (Anthropic) and OpenAI Codex.[12] For emergency managers and others who support EM functions, such products could, in principle, support planning, documentation, data integration, software maintenance, operational workflows, and analysis. However, the adoption of agentic AI products introduces additional governance, audit-

[10] National Institute of Standards and Technology (NIST), *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1, U.S. Department of Commerce, July 2024; Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung, "Survey of Hallucination in Natural Language Generation," *ACM Computing Surveys*, Vol. 55, No. 12, March 2023.

[11] Luis Aranda and Kasumi Sugimoto, "The Agentic AI Landscape and Its Conceptual Foundations," Organisation for Economic Co-operation and Development, OECD Artificial Intelligence Papers No. 56, February 2026.

[12] Anthropic, "Claude Code by Anthropic," webpage, undated; OpenAI, "Introducing Codex," webpage, May 16, 2025.

ability, and human-oversight concerns because their use often reduces transparency in how outputs are produced.[13] Effective adoption may therefore require a level of technical capacity—including the ability to evaluate intermediate steps and verify outputs.

Guidance from NIST identifies additional considerations for generative AI systems, including privacy, information integrity, transparency, third-party dependency, and governance controls.[14] Accordingly, many organizations and government agencies have established policies that restrict, govern, or require authorization for the entry of sensitive, classified, controlled, proprietary, or personally identifiable information (PII) into external AI systems unless such information is appropriately secured and the action is expressly authorized.[15]

Many Products Support Specific Tasks Within an EM Task Area but Would Require Stacking to Best Meet EM Needs

Across preparedness, response, and recovery, most commercially available AI products represent partial solutions relative to the task areas that they touch. *Rather than supporting the full scope of AI-amenable work within an EM task area, individual products typically address narrow slices of that work.* As a result, emergency managers and others who support EM functions and who seek coverage for the full scope of a task area would often need to stack multiple products together, and stacking could carry substantial cost, time, procurement, and integration burdens.

The partial-solutions pattern shows up in two related ways. First, most products we reviewed provide risk scores, hazard layers, detections, forecasts, transcripts, or drafts that could feed into an EM process, but they do not guide or facilitate the process. This is appropriate given that EM work is and should remain human-led. The issue is that the inputs the products supply are narrow enough that, even if emergency managers or others who support EM functions procure a stack of products, they would still have to spend additional effort assembling those inputs into a usable picture before they can concentrate on the work that the task area ultimately requires. This pattern surfaces across every functional area.

Second, *products fragment within task areas in ways that mean that no single product covers the full scope of AI-amenable work within a task area.* In some task areas, the fragmentation is sequential: The task area breaks down into stages, and products cluster in some stages while leaving others thinly covered. Search and Rescue, for example, has substantial AI product coverage for person detection via drone video, thermal imaging, and autonomous navigation in Global Positioning System (GPS)–denied environments but little coverage for the stabilize stage that connects rescue to medical triage. Coordination across multiple partners within a task area is rarely supported by AI-enabled products.

Most products provide single-domain feeds when the task area calls for cross-domain synthesis and application. Situational awareness for response (Situational Awareness, Hazard Monitoring, and Information Sharing) and recovery (Data, Information Management, and Situational Awareness for Recovery) is the clearest example of this phenomenon. This product area has the highest density, yet most products deliver only one kind of feed, such as weather, satellite imagery, social media, or IoT sensor data. Based on the available evi-

[13] Yonadav Shavit, Sandhini Agarwal, Miles Brundage, Steven Adler, Cullen O'Keefe, Rosie Campbell, Teddy Lee, Pamela Mishkin, Tyna Eloundou, Alan Hickey, et al., "Practices for Governing Agentic AI Systems," OpenAI, 2023.

[14] NIST, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, U.S. Department of Commerce, January 2023; NIST, 2024.

[15] Shalanda D. Young, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," memorandum for the heads of executive departments and agencies, Office of Management and Budget, M-24-10, March 28, 2024; NIST, 2023.

dence, we conclude that emergency managers and others who support EM functions who use one of these platforms for situational awareness would still need separate products for the multiple other data streams that they need to monitor during both response and recovery. Additionally, they would need to integrate, synthesize, and apply that data to inform decisions and operations.

Further compounding the stacking challenge, *many products in the landscape were not built for EM*. They originate in adjacent commercial sectors, such as insurance, reinsurance, commercial fleet management, enterprise IT, supply chain and logistics, and general-purpose contact centers, and were (or could be) adapted for EM use rather than designed for it. These products could support EM indirectly when others who support EM functions procure them to support day-to-day work and use them during disaster response and recovery. Yet, based on our analysis, many products on the market would require additional development to support EM explicitly. This pattern is visible across functional areas.

The practical consequence of these patterns is that stacking is not a neutral work-around. If emergency managers and others who support EM functions could procure and assemble multiple products to approximate the AI support they want for a given task area, doing so would introduce costs on several fronts. Procurement processes would need to be run for each product, often with technology companies, whose terms, pricing models, and contracting pathways differ considerably from each other. Products frequently do not integrate with one another in meaningful ways, which would require EM offices to move or reconcile data across products before the products become useful. Staff time would be required to learn and maintain multiple systems, and licensing costs would accumulate across products that each address only a slice of the work. *This dynamic also suggests an opportunity and innovation space for technology companies to explore and potentially remediate.*

A small number of products in the landscape are designed primarily as integration layers—platforms whose core function is connecting, aggregating, or orchestrating disparate systems rather than performing a discrete EM task. Such products as Overwatch Fusion+, Vantiq, and Microsoft Power Platform illustrate an emerging category of products that are AI-enabled and may reduce the burden of stacking by serving as connective infrastructure across other products.

There Are Clear Gaps in the Product Landscape

In the previous section, we described how products fragment within task areas such that no single product covers the full scope of AI-amenable work and why that pattern would require emergency managers and others who support EM functions to stack multiple products together to achieve meaningful coverage. In this section, we consider gaps of a different kind. Rather than looking at how coverage is divided within task areas, we look at where coverage is absent or sparse across the landscape. The gaps we describe are areas in which the market has produced comparatively fewer products, certain kinds of work are consistently underserved across multiple task areas, or the nature of the work has kept AI at the margins.

Multi-organization coordination is weakly automated. Most task areas involve coordination across multiple organizations working with the same survivor, event, or operation, and across these task areas, AI coverage to support such coordination is thin. For example, multiagency case management—orchestrating support around a single survivor across FEMA, the Small Business Administration (SBA), nonprofits, and state programs—remains weakly automated. Most case management platforms track a single organization's interactions with a survivor rather than coordinate across the full network of providers serving that household. We only found one product that supports interorganizational coordination among the organizations that coordinate volunteers and donations.

Preparedness task areas are underserved by purpose-built AI products. As we noted earlier, the product market thinly covers preparedness. This gap is notable given what preparedness work involves. The pre-

paredness task area framework used for this assessment establishes that hazard identification, vulnerability assessment, risk assessment, and integrated preparedness planning, among others (see Appendix E), each require input from a variety of players, the management of substantial data and documentation, and the synthesis and application of information across sources. Emergency managers spend most of their time on preparedness work.[16] The AI products that are available to support this work tend to be general solutions, such as enterprise document processing platforms, meeting and communication tools, and enterprise AI assistants, that technology companies have not yet adjusted to the specific workflows, terminology, or data needs of preparedness. Purpose-built preparedness products exist but are limited in number. Even when AI-enabled hazard and risk products do support preparedness work, they tend to stop at providing outputs: The outputs are not accompanied by AI support for translating them in order to inform decisions about what capabilities to build, where to invest in mitigation, or how to prioritize preparedness efforts remain largely manual.

With respect to recovery, products that support the actual decisionmaking, coordination, or execution of the work are not common. For example, disaster-specific behavioral health services at scale, including psychological first aid, grief counseling, and long-term trauma treatment for suddenly affected populations, have no dedicated AI products.

Product coverage varies substantially across task areas. Some task areas are thickly populated with AI products while others have very thin coverage. The thinnest areas include Emergency Debris Clearance and Route Opening, Fatality Management Services, and Capacity and Capability Analysis.

Survivor-facing work has thinner coverage than operational and technical work. Task areas that involve direct interaction with disaster survivors (e.g., Individual and Household Assistance and Case Management; Shelter, Mass Care, and Human Services; and Family Reunification and Welfare Checks) have substantially fewer AI products than operational or technical task areas. Where survivor-facing AI does exist, it is concentrated in chatbots and benefits navigation products, which help survivors find and apply for assistance programs. However other survivor-facing work has minimal AI support (e.g., the physical shelter operations of feeding, sanitation, and supply distribution). As with the product limitations we discussed in the previous section, these gaps are opportunities that technology companies can take advantage of.

Additional gaps that surfaced in the analysis do not fit neatly into any of these themes. These additional gaps at the task area level are cataloged in Appendix F.

Some Products Offer Novel Capabilities That Address EM Needs

We identified an array of products that do something many other products in the landscape do not (i.e., they offer relatively novel capabilities). Several categories of promising approaches are worth describing.

Across the landscape, some AI-enabled products expand participation and reduce coordination burdens by lowering the cost and reducing the complexity of training, planning, and operations. Approaches include immersive simulation platforms that generate dynamic, scenario-based training environments; products that are purpose-built for community-based organizations and formalize otherwise fragmented local capacity; and integrated platforms, such as digital twins and end-to-end preparedness systems that span multiple task areas, connect with existing systems, and reduce the need to stack multiple products. Together, these approaches point to a shift toward more accessible, comprehensive, and interoperable solutions that better align with the distributed nature of EM.

[16] National Preparedness Analytics Center, 2025.

If these products are commercially available, they are not as mature as other commercially deployed products. Some are in early production or pilot phases. These products show significant potential to support critical EM task areas, but they require additional time, development, or investment to achieve that potential. These products may currently focus on a narrow geographic region, a single hazard, or only part of a task area. There are far more products of this sort in the landscape than we describe here, and we do not suggest that the products we describe hold any more promise than others.

Conclusion

The landscape we characterized in this chapter is substantial and offers emergency managers and others who support EM functions real options for many of the task areas they work on. The depth of coverage varies depending on the task area, and the experience of navigating the landscape will differ for different buyers. While it is not yet known how well-suited LLMs are to EM tasks, these general-purpose products may be useful to emergency managers and others who support EM functions because of their low cost and ease of use.

We also uncovered reasons for optimism about where the landscape could go. The shape of the current landscape points to meaningful room for technology companies to develop products that address more of the work and for funders to direct investment toward the task areas and populations where coverage is thinnest. Beyond the purpose-built category, the substantial share of general solutions in the landscape signals that capable technology already exists in adjacent sectors and, with appropriate configuration and adaptation, could help address EM needs that purpose-built products do not yet reach.

The existence of a substantial landscape does not by itself answer whether emergency managers and others who support EM functions can put these products to use. Questions about what a product is, what it takes to stand a product up and maintain it, what a product costs, and what risks a product carries all sit between availability and adoption. In Chapter 3, we take up those questions directly.

CHAPTER 3

Procurement Realities: From Available to Adopted

A product that is on the market is not the equivalent of a product that is in use. In Chapter 2, we documented a broad and growing landscape of AI-enabled products with potential applications across preparedness, response, and recovery. Whether those products are positioned for adoption by the SLTT governments and nonprofit organizations working alongside them is a separate question and one that we take up directly in this chapter.

The gap between product availability and adoption is shaped by the structural features of what a product is, the technical and operational work required to stand up a product and keep it running, what a product costs, and what risks a product carries. Across each of these dimensions, the reality of moving a product from a technology company's website to working in an organization or jurisdiction involves complexities that public marketing does not reveal. These are complexities that emergency managers and others who support EM functions have to do substantial work to understand and compound in ways that the procurement conversation does not always anticipate. In this chapter, we walk through each dimension.[1]

One in Five Products in This Landscape Is a Component That Depends on Another Product to Function

Sixty-nine percent of the products in the landscape are stand-alone systems, sold as a single product with all functionalities included. An additional 13 percent are base products that anchor further paid capabilities, and 19 percent are those further capabilities, marketed and counted as products but dependent on a base product to function.[2]

This structure has a practical consequence for anyone evaluating what a product is. *Roughly one in five products in this landscape is a component, and the question of what else must be purchased for that product to work is a threshold question to answer before any other evaluation proceeds.* The ratio of add-on products to base products is greater than one, meaning that technology companies are marketing more extensions than

[1] We want to highlight a few things about what the information in this chapter can and cannot convey. First, the extent to which the findings we presented matter and whether they pose a barrier to potential procurement will vary by organization. We cannot state that the findings are inherently positive or negative in an organizational context for that reason, but we can and do suggest questions that emergency managers and others who support EM functions can ask technology companies when a finding has relevance for them. Second, these findings were assembled from public sources at scale: technology company–facing materials, procurement documents, cloud marketplaces, government records, and structured research across 1,179 products. Some findings rest on robust evidence across most of the landscape; others rest on evidence of varying quality, and where that is the case, we have tried to be explicit about what we could and could not determine. Pricing information was especially difficult to assemble, and we are candid about why. The reader should treat what follows as the best picture we could build from public evidence at scale and not as a complete account of every product in the landscape.

[2] These percentages do not sum to 100 percent because of rounding.

foundations. It also means that the path from an initial purchase to a full working configuration is often not a single transaction.

If a customer does not consider this interdependency prior to their initial purchase, they risk either inadvertently locking themselves into a particular brand or overlooking these unforeseen costs for future purchases. Either way, emergency managers should maintain awareness of these inter-product dependencies early on in procurement decisionmaking to reduce future costs or interrupted workflows.

Product dependency has several forms in practice, and these forms carry different weights for emergency managers and others who support EM functions. Some add-on products are features that are activated by changing a setting in a platform an organization already has. Others are separately licensed modules that require paying for and maintaining the base product alongside the add-on, with two contracts, two renewal cycles, and two sets of life-cycle considerations to manage. Some are capabilities that become available through a subscription tier upgrade rather than a separate purchase. Other add-ons require the technology company to activate them on the customer's behalf after purchase. And some are sold within a technology company's own ecosystem, in which the add-on only functions as part of a broader platform the organization must have already procured. *Emergency managers and others who support EM functions encountering any of these add-ons are not making a single decision about a single product. They are also making a decision about a relationship with a parent product, and the two decisions cannot be cleanly separated.*

A compounding factor is that *a product's marketing does not always make dependency obvious*. Capabilities are routinely listed, promoted, and sold as stand-alone products, even when they structurally require a parent product to function. It is worth noting that the dependency question is distinct from the question of how much effort is required to stand a product up. A component can be straightforward to activate within its parent platform, such as an administrator toggling on a setting, or it can require substantial technical work to configure and connect, such as integrating a licensed module into an existing enterprise system. The burden of standing up a product is a separate consideration, which we address later in this chapter. *Emergency managers and others who support EM functions may want to discern up front whether the product they are looking at is a complete working product on its own or a component that requires something else to function.*

For an emergency manager or EM-supporting organization, the implication extends beyond the initial product purchase. Buying an add-on is a commitment to the parent product's continued licensing, support, and life cycle for as long as the add-on is in use. If the parent product is depreciated, repriced, or restructured, the add-on is exposed to those changes regardless of how well it has served the organization. *A useful question to ask before procurement is not only what else must be purchased for this product to work, but what ongoing relationship with the technology company the organization is willing and able to sustain. The opportunity for technology companies is to develop clear information about product dependency and products and technologies that will alleviate this challenge.*

Most Products Are Delivered as Cloud Services, and Many Offer Additional Deployment Options

Nearly nine in ten products in the landscape are delivered as software as a service (SaaS). In a SaaS product, the software runs on the technology company's computers, and the customer accesses it through a web browser or an app. Most cloud-based products work this way. SaaS applies to 88 percent of products in this landscape. A small share of products classified as SaaS are free rather than subscription based.

Beyond SaaS, the other deployment categories each apply to less than 20 percent of the landscape, including on-premises (on-prem) capability,[3] hardware,[4] locally installed software, and edge.[5] Forty-six percent of products are offered in two or more deployment models. A drone is usually both hardware and edge. An enterprise platform may be offered as both SaaS and on-prem. A cloud service may also have a mobile app.

Among the 16 percent of products that include a hardware component, pure hardware is unusual. Most hardware products come bundled with software, and that software is most often cloud-delivered. A drone, a wildfire camera, or a field sensor is typically sold along with a cloud platform that handles analytics, fleet management, or data storage, and the cloud platform usually comes with ongoing subscription fees on top of the hardware purchase. A smaller group of hardware products (e.g., emergency services platforms, autonomous drones with onboard intelligence, networking appliances with cloud management) are built so that the hardware and software are genuinely inseparable. These products create a specific procurement situation: The hardware and software are purchased together and cannot be separated later, which means that emergency managers and others who support EM functions who want to switch one layer must switch both. Products in this category tend to be the most complex to procure, deploy, and maintain.

To benefit from most products, a jurisdiction's IT environment needs to be able to support cloud access. Hardware purchases almost always come with ongoing software costs, and those costs should be factored into any comparison between a hardware product and a software-only alternative. Purchasing integrated hardware and software products (the kind where one cannot function without the other) commits a jurisdiction to both layers from the same technology company, so the procurement decision is effectively a long-term commitment. On-prem is available for a meaningful share of products, but it is usually offered as an alternative to the technology company's cloud version rather than the main way the product is delivered. Emergency managers and others who support EM functions who require on-prem should confirm directly with the technology company that it is fully supported rather than a legacy option. Finally, because so many products carry multiple deployment labels, a product may have capabilities, such as mobile apps, offline modes, API access, or alternative hosting, that are not obvious from its primary marketing and are worth asking about.

The Path from Purchase to a Working, Sustained Product Is Often Neither a Single Transaction nor a One-Time Setup

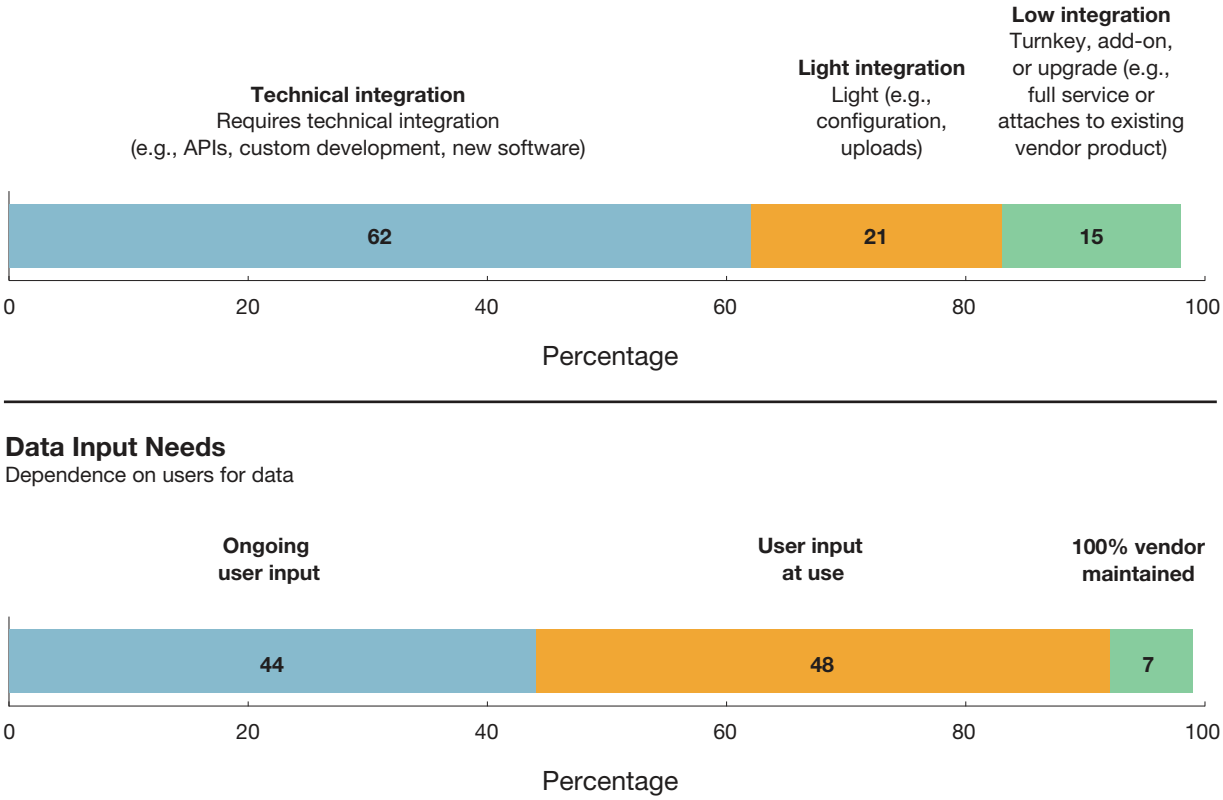
Most products in this landscape require meaningful effort to stand up and meaningful effort to keep running. Sixty-two percent of products require technical integration, meaning that technical expertise is required to connect the product to the systems an organization already uses (Figure 3.1). In most cases, that work falls on the customer's IT team or an outside integrator; in some cases, the technology company performs the work for the customer, but the work itself is still technical in nature and takes time to complete. Another 21 percent of products require what we call *light integration*: configuration, data entry, or a guided setup that a nontechnical administrator can typically handle without writing code or calling for IT support. Only 15 percent of products function as *turnkey*—usable immediately with no integration work required by the customer. For the remaining small share of products, we could not determine the integration burden from available evidence.

[3] *On-premises capability* means that the product can be installed and run on the customer's own servers or infrastructure rather than those of the technology company.

[4] *Hardware* means that the technology company ships a physical device, such as a drone, sensor, or camera.

[5] *Edge* is a term that describes where processing happens rather than a type of product. In practice, edge refers to processing that occurs on a device in the field or on equipment at a customer's site rather than in a centralized cloud.

FIGURE 3.1
Product Burden



NOTE: The burden levels reflect the degree of technical demand placed on the customer. One to two percent of products were *unknown* on each dimension of burden. Percentages may not sum to 100 percent because of rounding. The use of color and shading has no inherent meaning. User input may or may not be in combination with technology company–maintained data.

Once products are deployed, most need continuous attention to operate. Forty-four percent of products require ongoing data input from users, meaning that staff must maintain access to data (e.g., entering records, updating contact lists, streaming sensors or camera data). Another 48 percent require user input at the time of use, meaning that the product needs specific files, documents, prompts, or location data each time the user engages with the product. In both cases, technology companies often supply data of their own alongside what users provide. For example, a mass notification system may come with geographic or demographic data that the user's contact list layers onto. User input is rarely the only data that a product runs on, but it is data that the product cannot achieve full functionality without. Only 7 percent of products are fully technology company–maintained, arriving with the data they need already in place, and they are not reliant on the customer to keep them supplied. Table 3.1 shows the kinds of data that are associated with each input and common product types for each input.

Altogether, these patterns suggest that a typical product in this landscape requires technical expertise and work to deploy, configuration or integration to connect it to existing systems, and continued data inputs to operate. These findings carry direct consequences for any emergency manager or EM-supporting organiza-

TABLE 3.1
Examples of User-Required Data Inputs and Associated Product Types

Required Data Input	Examples of Data Types That Are Associated with Products Requiring User Input at Use	Examples of Associated Product Types
User input at use	• Audio, speech, and video files • Documents (e.g., pdfs, Word documents, Excel spreadsheets, comma-separated values files) • Records (e.g., bank statements, receipts, invoices) • Natural language text (e.g., prompts, queries, search terms, message content) • Geospatial inputs (e.g., addresses, coordinates, areas of interest drawn as polygons, shapefiles, keyhole markup language and keyhole markup zipped files)	• Damage inspection products • Drone operation • Document processing • Translation services • Geospatial intelligence • Satellite tasking
Ongoing user input	• Staff continuously enters records, cases, incidents, or transactions through daily use. • User-deployed sensors or devices continuously stream data into the platform. • The user continuously updates contact lists, groups, and recipient data. • User cameras provide real-time video feeds. • The user inputs incident and situation data. • The user provides ongoing platform access to databases.	• Incident management • Case management • Dispatch and CAD • Records management • Workforce and volunteer management • Resource management and supply chain • Real-time communications platforms

tion who is weighing a purchase. Technical integration draws on IT capacity that not every EM office has in house, and hiring or contracting for such capacity adds cost and time that the product's listed price does not reflect. Products that depend on ongoing data input add operational load that scales with how heavily the product is used; that load is not likely to ease during periods when the product matters most. *A useful question to ask before procurement is not only whether a product meets a functional need but whether the organization has the capacity to deploy the product properly and sustain it under the conditions it is meant to support.*

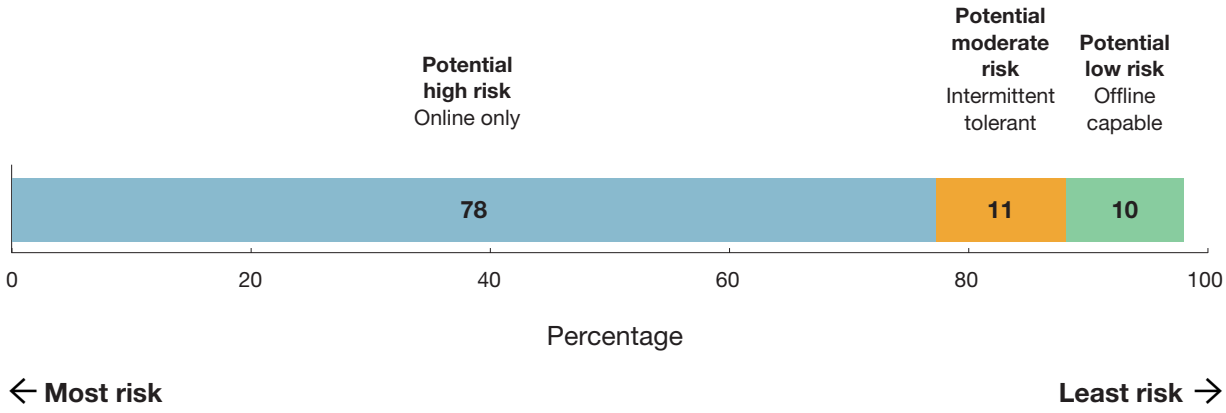
Most Products Stop Working When Connectivity Fails

Connectivity refers to whether a product always functions under the kinds of conditions EM work can impose, particularly response. Approximately 78 percent of products require a continuous online connection to function (Figure 3.2). If the connection drops, the product stops working. Another 11 percent of products can tolerate intermittent connectivity, meaning that they can continue to function during brief disruptions and reconnect when service returns. Only 10 percent of products can function fully offline.

For an emergency manager or EM-supporting organization, this finding has two practical points. First, connectivity often fails under disaster conditions. A product that requires continuous connection is a product that may be unavailable during the exact response circumstances it was procured to support. Buyers who are evaluating such products should understand the specific failure modes of those products: What happens if the connection drops for ten minutes during an active incident, and what happens if it drops for a day or more?

Second, a product's connectivity dependence is not only about a jurisdiction's internet access. Most cloud-hosted products also depend on the technology company's server infrastructure remaining available. A product may be reachable from a fully functional local network and still be unusable because the technology company's servers are down, overloaded, or affected by a regional event. This second dependency is rarely visible at the procurement stage. *It may be worth asking technology companies directly about their own continuity commitments: What uptime do they guarantee? Is their infrastructure geographically distributed? What are the customer's options if the technology company's systems are unavailable during an incident?*

FIGURE 3.2
Potential Connectivity Risk



NOTE: *Unknown* products were 0.2 percent. The use of color and shading has no inherent meaning. Percentages may not sum to 100 percent because of rounding.

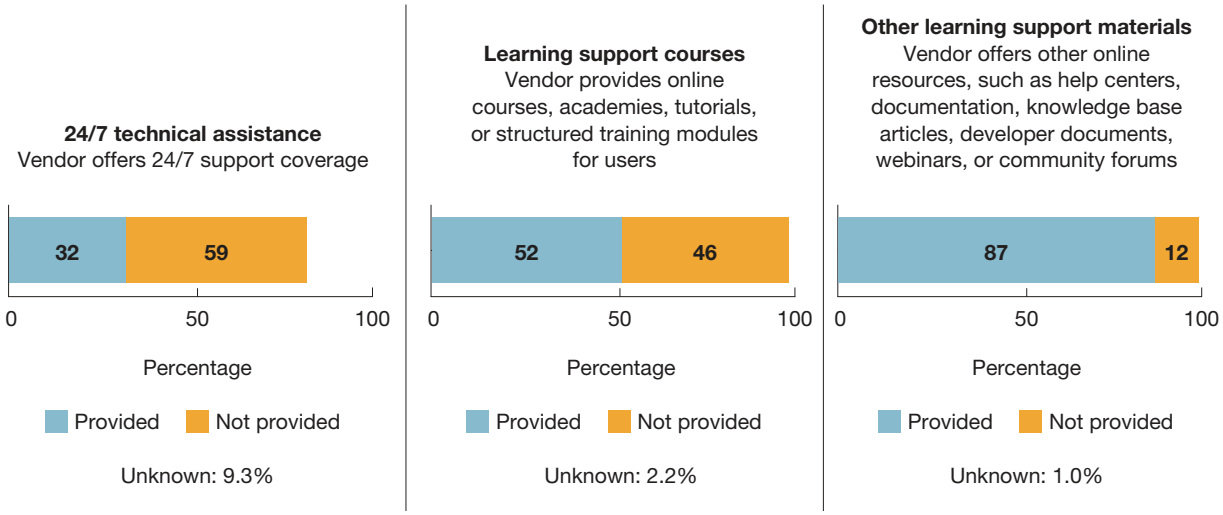
Most Technology Companies Provide Learning Support for Users but Not 24/7 Technical Assistance

Sixty-eight percent of technology companies do not publicly advertise 24/7 technical assistance (Figure 3.3). A little more than half offer structured learning courses, and 87 percent offer some form of supplementary learning materials, such as help centers, documentation, or community forums. A small share of technology companies state that training is not required to use their product. Available support does not vary significantly across preparedness, response, and recovery products.

The depth of support that a technology company offers tends to track the product's deployment scale. Products at an earlier stage of deployment or with narrower deployment tend to come with basic documentation. Products with broader deployment more often come with structured documentation and webinars, and the most broadly deployed products are the most likely to come with structured training programs, comprehensive documentation, and 24/7 support.

For any EM buyer procuring a product that is meant to function during an active response, the 24/7 finding is the one that matters the most. Disasters do not follow business hours, and a product failure during a response operation may require immediate technology company assistance to resolve.

FIGURE 3.3
Provision of Technical Assistance and Learning Support



NOTE: The use of color and shading has no inherent meaning. Percentages may not sum to 100 percent because of rounding.

Pricing Is Largely Opaque, and Access Models Vary Widely

How a product is priced in this landscape varies substantially, and how much a product costs is rarely answerable from public sources. These two findings describe a procurement reality that compounds the integration, support, and composition complexities addressed earlier in this chapter.

We attempted to estimate a primary pricing basis for most products in the landscape.[6] From what we could discern, four structural pricing approaches appear to account for the products in that group. Approximately half of those products (49 percent) are priced per organization through a flat-fee or enterprise-license arrangement. Seventeen percent are priced per user, scaling with the number of individuals who are accessing the product. Sixteen percent are priced per device, scaling with the number of physical units, such as cameras, drones, sensors, or tablets, that are deployed. Fifteen percent are priced on a usage basis, and charges accrue with such activity as API calls, processed pages, minutes of audio, or messages sent. The pricing approach for the remaining 4 percent of products could not be determined.

These four categories capture the dominant pricing structures, but they are an oversimplification of the underlying variation. The pricing mechanisms we documented across the landscape include the following:

- one-time arrangements—the outright purchase, patent licensing, and technology transfer licensing for government-developed products
- recurring arrangements—subscription and per-period licensing
- per-unit scaling—per seat (by user, device, facility, or property), per use, per unit, per image, per event, per character, per minute, per token, per API call, per page processed, and per message
- quote-based or configuration-dependent—sales-led custom-quote pricing and tier-dependent pricing that varies with selected configuration.

[6] We considered 1,134 products (96 percent) potentially estimable for pricing and only attempted to estimate pricing for those products.

The category to which we initially assigned a product is a starting description rather than a complete one. The granular pricing mechanism that emergency managers and others who support EM functions will encounter (e.g., per seat by property versus by user, per token versus per API call, tier-dependent or configuration-dependent) is often more specific than the category labels we used would suggest. Two additional realities compound this granularity. First, approximately one in five products in the landscape is an add-on that requires a base product to function, and the add-on and base product could carry different primary pricing bases. A hardware product that is purchased outright often requires a separate subscription for the software that maximizes its utility, and that subscription could be priced per user, per device, or per use depending on the technology company. Second, most products will require integration with the systems that a jurisdiction already uses or will need to procure to use the product; any new systems that a jurisdiction will need to procure may carry pricing bases that are different from the primary product. The total cost for a working deployment would often be a composite of multiple pricing bases, and emergency managers and others who support EM functions cannot project the full cost from the primary product's access model alone.

Determining what a product costs required substantial research effort using various strategies, and even such substantial effort did not produce evidence of overall quality that we could be confident in.[7] Most technology companies in this landscape do not publish pricing on their own websites; most do not list products in cloud marketplaces, and most are not priced in federal procurement systems. When we could find prices, they were most often buried in local government meeting documents or cooperative contract schedules that reflected specific negotiated arrangements for specific jurisdictions at specific points in time. These sources do not establish list prices, and the prices we did find are not necessarily reproducible for a different buyer with a different profile. For products for which none of these sources yielded usable evidence, our research relied on market analogy estimates drawn from comparable products, an approach that is inherently uncertain. The effort required to assemble even partial pricing evidence across this landscape may exceed the time and effort an emergency manager or others supporting EM functions can devote to a single purchasing decision. *Useful questions to ask before procurement may be not only what the primary product costs, but what else must be purchased for it to function, what pricing basis applies to each component, and what the total first-year and ongoing costs would look like for deployment at the organization's scale.*

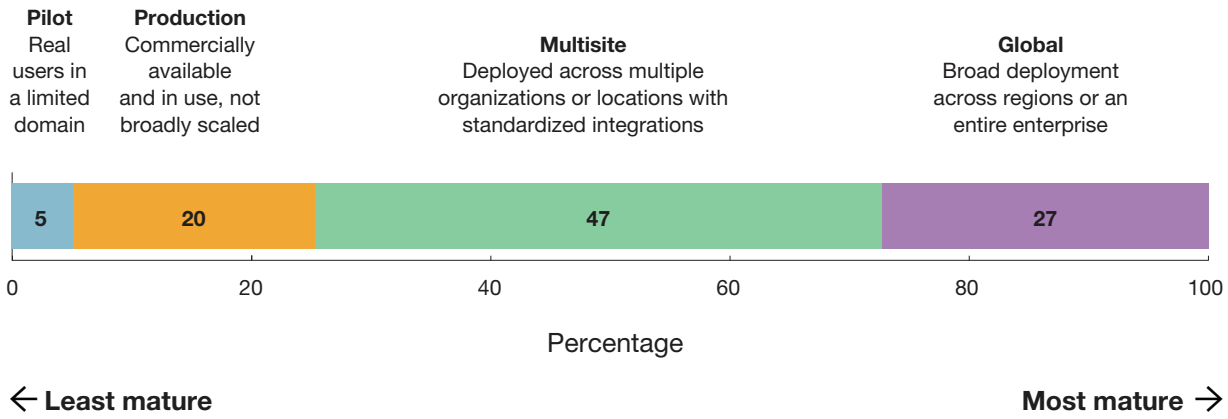
Most Products Have Documented Deployment Across Multiple Sites or Countries

Nearly three-quarters of products in the landscape have documented deployment across multiple sites or countries, and only 5 percent remain at pilot stage (Figure 3.4). Forty-seven percent of products are deployed at multiple sites, varying from three sites to thousands of sites. Twenty-seven percent are deployed globally, with documented adoption across multiple countries and/or continents. The remaining 20 percent of products are commercially available and in supported use but without publicly documented evidence of broad multi-organizational deployment. Pilot-stage activity tracks overall market distribution across functional areas.[8]

[7] We pursued pricing through a systematic ladder of sources: technology company pricing pages, Amazon Web Services and Azure marketplaces, federal procurement systems (e.g., GSA Advantage, USAspending, SAM.gov), United Kingdom Digital Marketplace listings as foreign exchange–converted proxies, city council agendas, county board resolutions, state sole-source memorandums, federal award databases, and cooperative contract price schedules from such vehicles as TXShare, OMNIA Partners, HGACBuy, Sourcewell, state departments of information resources contracts, and the National Cooperative Purchasing Alliance.

[8] Products at the pilot stage may be less likely to have an online presence. Therefore, there may be more products to support EM functions at the pilot stage than the available evidence suggests.

FIGURE 3.4
Product Deployment Scale



NOTE: The use of color and shading has no inherent meaning. *Unknown* products were 0.3 percent. Percentages may not sum to 100 percent because of rounding.

Multisite and global deployment demonstrate different things. Multisite products have documented use across multiple customer sites, and the examples in this category (e.g., Juvare products, Skydio products, Carbyne products, AiDASH, Urbint products, Generac ARC) are dominated by EM purpose-built response products and products that support task areas that utilities and infrastructure owner-operators manage. Global deployment, in contrast, is concentrated among hyperscalers and enterprise platforms (e.g., Google Cloud, Microsoft Azure, Amazon Web Services, Salesforce, Palantir) and in broadly adaptable administrative products (e.g., Zoom AI Companion, Copilot, Zendesk Suite).[9] These products are demonstrably used at scale, but that adoption is not necessarily by organizations that support EM. Public evidence of deployment is not evidence of performance, usability, or long-term sustainment by entities that implement the products. Twenty percent of products categorized as *in production* are commercially available and supported. We could not find evidence of broad deployment of these products across organizations.

Several things apply for emergency managers and others who support EM functions who are evaluating a product in this landscape. The products surfaced here have generally moved past the pilot stage, which is a useful baseline. Beyond that, the deployment category that a product falls into answers a specific question and not a broader one. A product that falls into in the multisite category has documented use at multiple EM or EM-adjacent customer sites, and emergency managers and others who support EM functions can reasonably ask which sites and contact those sites. A product that falls into the global category has documented use at scale, but whether that scale includes EM use is a separate inquiry. A product that falls into the production category is in supported use somewhere, but the buyer should expect to ask the technology company directly for deployment counts, customer references, and named installations rather than infer that information from public sources.

[9] *Hyperscaler* means a company that operates cloud infrastructure at massive global scale, with enormous data center capacity and the ability to rapidly expand computing, storage, and networking services on demand (e.g., Amazon Web Services, Google Cloud).

The Type and Degree of Risk Vary Across Products

Nearly every product in the landscape reflects some form of risk, although the prevalence and character of that risk vary considerably across the dimensions we examined. Privacy risk is near-universal (Figure 3.5). Legal and cyber concerns apply to a minority of products. AI governance assurance is rare across the landscape. Each dimension carries different procurement implications and is addressed separately in the following sections.

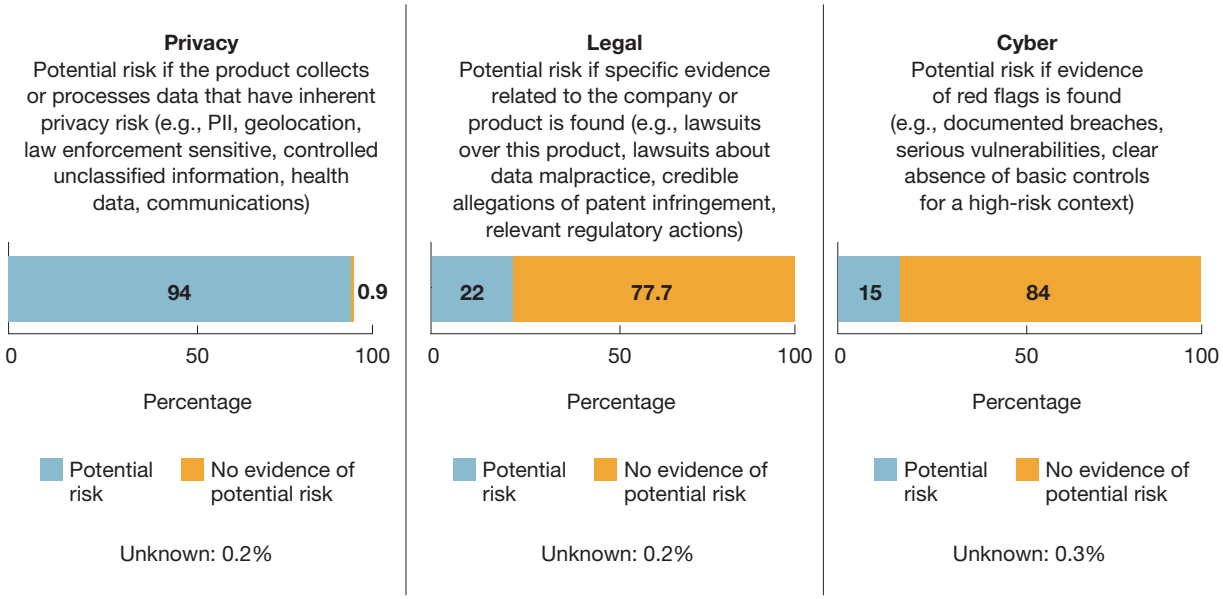
Potential Privacy Risk Is Near-Universal, and Most Products Trigger Multiple Categories at Once

We identified potential privacy risk for 94 percent of products. We considered a product to present privacy risk if we found evidence that it collects or processes data that are inherently privacy sensitive, including PII, geolocation, health data, law enforcement–sensitive data, controlled unclassified information, or communications content. Some technology companies have taken steps to mitigate the risks associated with collecting and maintaining private data, such as completing security certifications. Regardless, the overwhelming majority of products require the use of some type of privacy-sensitive data.

Geolocation is the most common privacy-sensitive data type in the landscape (Table 3.2). PII is the next most common form of data that products require.

Most products do not trigger a single privacy category; they typically trigger three, four, or five or more categories simultaneously. An emergency services or public safety answering point system routinely combines PII, geolocation, communications content, health data, and law enforcement–sensitive data within a single call event. A drone-as-first-responder platform combines aerial imagery, geolocation, video surveillance, law enforcement–sensitive data, and audio recordings. An emergency medical services documentation system combines protected health information, PII, geolocation, communications, and financial and billing

FIGURE 3.5
Potential Risk from Evidence of Privacy, Legal, or Cyber Concerns



NOTE: The use of color and shading has no inherent meaning. Percentages may not sum to 100 percent because of rounding.

TABLE 3.2
Most Observed Patterns in the Privacy Concerns Data

Rank	Pattern Label	Pattern Description
1	Geospatial and location-linked data	• Real-time GPS tracking of individuals, vehicles, or assets • Property- and address-level risk scores and aerial imagery • Population-level mobility analytics • Critical infrastructure asset geolocation
2	Visual surveillance, imagery, and biometric data	• Drone footage, security camera feeds, body-worn camera recordings, and satellite imagery capturing people or private property • Biometric identifiers, including facial recognition templates, voice prints, DNA and genomic profiles, and fingerprints
3	Public safety, health, and emergency communications data	• Protected health information and patient records • EMS care reports and behavioral health data • Recorded or transcribed emergency services calls • Dispatch and CAD incident data • Criminal justice information services–regulated criminal justice records • Crisis hotline content
4	Financial, insurance, and workforce monitoring data	• Loan applications and credit records • Policyholder claims and underwriting inputs • Donor and payment card data • Continuous driver behavior tracking, in-cab video, wearable physiological sensors, and GPS-verified timesheets tied to named individuals
5	Digital intelligence and third-party AI data flows	• Aggregation of social media content and geolocation-tagged posts for OSINT • Dark web credential and breach monitoring • User prompts, organizational documents, or PII routed to external LLM providers as part of generative AI features

NOTE: General PII collected through routine platform operations (e.g., account registration, contact details, internet protocol addresses) is present in virtually every product flagged "Yes" and omitted from this table because it does not differentiate among products. EMS = emergency medical services.

data. This stacking means that the privacy implications of a single product often involve multiple privacy categories at once, which compounds the exposure beyond the risk of any single category.

For the few products for which we found no evidence of potential privacy concerns, their core data pipelines involve no PII, no individual geolocation, no health data, no communications content, no biometric data, and no law enforcement–sensitive data. *For an emergency manager or EM-supporting organization, this means that avoiding privacy considerations is not really an option within this landscape. The universe of AI products that do not raise any privacy questions is very small, and procurement decisions therefore require engaging with privacy factors in nearly every adoption decision process.*

Potential Legal Concerns Were Identified for About One in Five Products

We identified potential legal concerns for 22 percent of products. We flagged a product for potential legal concerns when specific evidence pointed to lawsuits involving the technology company or product, allegations of data malpractice, credible patent infringement claims, or relevant regulatory actions.[10]

Five patterns emerged from the documented legal concerns, which we ranked from most to least common (Table 3.3). Patent infringement litigation appeared most often, naming technology companies or products as either plaintiffs or defendants in U.S. patent lawsuits. Data privacy and data malpractice class actions were

[10] The absence of public evidence of legal concerns, as we define it, is not proof that such concern does not exist with respect to the technology companies or specific products that were not flagged for potential legal risk.

TABLE 3.3
Most-Observed Patterns in the Potential Legal Concerns Data

Rank	Pattern Label	Pattern Description
1	Patent infringement litigation	The technology company or product is named in U.S. patent infringement lawsuits, either as a plaintiff or a defendant.
2	Data privacy and data malpractice class actions	Lawsuits allege the unauthorized collection, unauthorized sharing, or mishandling of personal data.
3	Regulatory enforcement actions	Government agencies (FTC, FCC, SEC, state AGs) are taking formal action against the technology company.
4	Securities fraud class actions	Shareholder lawsuits allege that the technology company made false or misleading statements about business performance.
5	Contract, commercial, or employment disputes	These are business-to-business disputes, employment lawsuits, or commercial litigation that do not fit other categories.

the second-most frequent, alleging the unauthorized collection, unauthorized sharing, or mishandling of personal data. This category also includes the growing body of lawsuits against major AI platform technology companies over training data practices. Regulatory enforcement actions were third, involving formal action by such agencies as the Federal Trade Commission (FTC), Federal Communications Commission (FCC), Securities and Exchange Commission (SEC), or state attorneys general (AGs). Securities fraud class actions were fourth, brought by shareholders alleging that technology companies made false or misleading statements about business performance. Contract, commercial, or employment disputes were fifth, covering business-to-business litigation, employment lawsuits, and commercial cases that do not fit the other categories. While patent litigation is common by volume, privacy and AI-related copyright cases more often impose direct operational risks, including regulatory fines, restrictions on data use, and requirements for product modification.[11]

Within the 22 percent of products flagged, most triggered more than one kind of legal concern. Flagged products typically had two or three overlapping concerns. The potential legal risks we identified are concentrated most heavily among products that support Situational Awareness, Hazard Monitoring, and Information Sharing (58 percent of legal concerns) and Incident Coordination and Operational Decision Support (32 percent of legal concerns), the task areas at the center of active response operations.

For an emergency manager or EM-supporting organization, the practical takeaway from this finding is that a documented legal concern means that something specific has been filed or pursued against the technology company or product. The nature of the concern varies, and not every legal action will have bearing on a jurisdiction's operational use of the product. Data privacy class actions provide the most direct line to how a product actually handles information during use because they contest the data practices that the product relies on. *Emergency managers and others who support EM functions who encounter a legal finding during their due diligence should ask what the specific concern is* and whether that concern pertains to the product being considered or a different product from the same technology company and whether any outcome has been reached.

[11] Administrative Office of the U.S. Courts, *Judicial Business of the United States Courts: 2023 Annual Report*, U.S. Courts, 2024; European Data Protection Board, *Annual Report 2023: Safeguarding Individuals' Digital Rights*, 2024; Pamela Samuelson, "Generative AI Meets Copyright," *Science*, Vol. 381, No. 6654, 2023.

Potential Cyber Concerns Were Identified for 15 Percent of Products

We identified potential cyber concerns for 15 percent of products (Figure 3.5). A product was flagged for cyber concerns when specific evidence pointed to documented breaches, serious vulnerabilities, or a clear absence of basic security controls in high-risk contexts.[12] General uncertainty about a technology company's security posture was not enough to trigger the flag.

Five patterns emerged from the documented cyber concerns, which we ranked from most to least common (Table 3.4). Documented breaches or ransomware attacks was the most frequently observed pattern.

Within the 15 percent of products flagged, most triggered more than one kind of cyber concern. Flagged products frequently had two and occasionally three overlapping concerns.

For an emergency manager or EM-supporting organization, three points from this finding are worth noting. First, certification alone is not a reliable signal of a product's current cyber posture. Technology companies with recognized certifications still appeared in the cyber concerns data because of confirmed breaches or unpatched critical vulnerabilities, which means that emergency managers and others who support EM functions cannot treat the presence of a certification as the end of the cyber inquiry. Second, AI-enabled products introduce attack surfaces that do not exist in traditional software, and technology companies and customers may have little familiarity with these new risks. Prompt injections, model manipulation, and guardrail bypass are specific to AI systems and are not addressed by traditional cybersecurity controls. Third, EM operations depend on information being accurate, available, and protected; the cyber concerns documented here apply to a minority of products but are meaningful to consider. Evaluating a product's cyber posture during procurement rather than after deployment is therefore a practical step.

TABLE 3.4
Most-Observed Patterns in the Potential Cyber Concerns Data

Rank	Pattern Label	Pattern Description
1	Documented breaches or ransomware attacks	The technology company or product has experienced a confirmed, publicly reported security incident, whether or not the technology company has earned one or more cyber certifications since the incident (e.g., SOC 2, ISO 27001).
2	Published critical vulnerabilities	These are products with formally documented, high-severity security vulnerabilities, especially those in the Cybersecurity and Infrastructure Security Agency's Known Exploited Vulnerabilities Catalog.
3	Inherent vulnerabilities in IoT and hardware devices, such as drones and cameras	These are connected devices (e.g., drones, cameras, sensors, wearables) with category-wide known vulnerabilities.
4	Inherent vulnerabilities in AI and ML systems	These are vulnerabilities that are unique to AI and ML systems that introduce novel AI-specific attack surfaces (e.g., prompt injection, model manipulation, guardrail bypass).
5	Sensitive data handling combined with documented misuse or privacy concerns	Products for which the nature of the data collected combined with documented misuse or questionable practices constitute a cyber concern.

NOTE: ISO = International Organization for Standardization; SOC = System and Organization Controls.

[12] The absence of public evidence of cyber concerns, as we define it, is not proof that such concern does not exist for products that were not flagged for potential cyber risk.

Verifiable AI Governance Claims Are Rare in the Product Landscape

Two governance standards anchor the current conversation about responsible AI: ISO and the International Electrotechnical Commission's (IEC's) ISO/IEC 42001 and NIST's AI Risk Management Framework.[13] ISO/IEC 42001, published in December 2023, is the first globally recognized certification for AI governance.[14] In the landscape, less than 1 percent of products carry a technology company claim to either standard.[15] For an emergency manager or EM-supporting organization, this is the operative finding. *Most products that are being marketed for EM use do not currently offer verifiable AI governance assurance* through the only internationally standardized mechanism available (ISO/IEC 42001), and claims of alignment with the U.S. voluntary framework cannot be independently checked. The implication is not that AI products should be avoided; the implication is that the assurance layer that exists for cybersecurity (e.g., SOC 2, ISO 27001, the Federal Risk and Authorization Management Program [FedRAMP]) has no equivalent in practical operation for AI governance today. *Emergency managers and others who support EM functions who are asking, "How do I know this product's AI is responsibly governed?" will find that the available answers depend mostly on what the technology company chooses to tell them*.

About Six in Ten Technology Companies Publicly List Evidence of Cyber Risk Mitigation

Approximately 60 percent of technology companies in the landscape publicly list evidence of cyber risk mitigation efforts, such as security certifications,[16] data privacy certifications,[17] or formal risk management documentation. The remaining 40 percent do not. The absence of publicly listed evidence is not proof that a certification does not exist.[18]

A technology company with no public certifications is not automatically untrustworthy, but emergency managers and others who support EM functions in that situation will need to rely more heavily on contractual language and direct technology company inquiry than on third-party attestation. A technology com-

---

[13] NIST, 2023; ISO and IEC, *Information Technology—Artificial Intelligence—Management System*, ISO/IEC 42001:2023, 1st ed., December 2023.

[14] ISO/IEC 42001 is, as of 2025, the only globally recognized standard that combines (1) an AI-specific scope, (2) a formal management system structure, and (3) a pathway to independent, accredited third-party certification of organizational AI governance practices. The American National Standards Institute's National Accreditation Board launched its accreditation program for ISO/IEC 42001 in January 2024. The NIST AI Risk Management Framework, published in January 2023, is the primary U.S. reference. It is a voluntary framework organized around four functions (govern, map, measure, and manage) and was augmented by a 2024 NIST profile specific to generative AI (NIST, 2024). Unlike ISO/IEC 42001, NIST's AI Risk Management Framework has no certification mechanism, meaning that an organization can claim alignment, but no third party can independently verify that claim.

[15] The absence of public evidence of compliance with an AI governance standard is not proof that technology companies do not have such evidence or that AI is ill-governed.

[16] Among technology companies that did publicly list evidence, certification acquisition followed a recognizable progression. SOC 2 Type II served as the foundation, followed by ISO 27001 and domain-specific certifications.

[17] Two certifications specifically address data privacy: ISO 27701 certifies privacy information management for PII, and ISO 27018 certifies PII protection specifically in public cloud environments. Both require independent third-party audits on multiyear cycles and are the most commonly cited privacy-specific certifications among technology companies in this landscape.

[18] For example, in the United States, data protection obligations are often imposed through technology company contracts and data-processing agreements rather than through certifications alone. U.S. privacy laws at both the state and federal levels require contractual provisions that govern the handling of personal data, and these agreements are legally enforceable through contract law and regulatory oversight. Thus, the absence of a certification does not necessarily mean the absence of data protection commitments.

pany with public certifications still warrants follow-up questions because the certifications confirm that certain processes have been audited but do not confirm that the technology company's practices meet the buyer's specific requirements. In either case, *the procurement conversation needs to go beyond the question, "Do you have certifications?" and into questions about what those certifications cover, what they do not cover, and what the technology company's contractual commitments look like.*

Conclusion

In this chapter, we relied on the best public evidence available at scale for our analysis, using product-facing materials, procurement records, marketplace listings, and structured research across 1,179 products. Those sources cannot capture every implementation detail or local purchasing reality, but they are sufficient to identify an important pattern: The central challenge is not simply whether AI-enabled products exist for EM but whether those products are practically procurable, deployable, and sustainable for the organizations that might use them.

Considered individually, the findings in this chapter highlight the manageable realities of procurement, including product dependencies, integration demands, connectivity requirements, opaque pricing, uneven support, and varied risk profiles. Altogether, the findings suggest that getting to the point of implementation, much less full adoption, depends on not only product availability but also the organizational capacity to navigate these complexities. Some EM offices and others who support EM functions will be better positioned to do this than others. Technology companies should seriously weigh these factors that will limit procurement and, therefore, adoption and diffusion. In Chapter 4, we examine the conditions that may shape adoption and diffusion across EM.

CHAPTER 4

From Awareness to Broad Adoption: What Will Shape AI Diffusion in Emergency Management?

In Chapters 2 and 3 we established, respectively, the landscape of AI-enabled products that are available to support EM and what working with those products may require. In this chapter, we address the report's two remaining research questions: whether the conditions for adoption and diffusion of AI products across EM appear to be in place and what the available evidence suggests may need to change for those processes to occur.

AI adoption that is relevant to EM can occur through two primary channels. The first is direct adoption within EM offices. The second is adoption by other organizations that support EM functions, such as transportation, utilities, health care, public safety, logistics, and communications. Because EM outcomes depend on both channels, we examine adoption across EM offices in greater depth while treating adoption across EM-supporting organizations more generally. We draw on findings from Chapters 2 and 3, along with relevant literature on AI and technology adoption in EM, SLTTs, and nonprofits, to identify the conditions that are likely to support or hinder adoption and diffusion.[1]

Adoption Within EM Offices Is Emerging at Different Rates

In Chapter 1, we described organizational technology, or product adoption, as a progression from awareness and evaluation to implementation, sustainment, and eventual standardization.[2] When organizations achieve standardization, the adoption of a product is complete. The available evidence suggests that many EM offices remain concentrated in the earlier stages of that process (e.g., awareness) for AI-enabled products rather than at meaningful implementation or beyond.

We draw this conclusion from the findings of a 2025 landmark study of EM offices across the country.[3] This study reports that such technologies as public alert and warning systems, GIS, social media platforms, and virtual EOCs are widely accessible across all SLTT levels.[4] AI is the outlier. It is the least accessed of the

[1] We rely heavily on the extant literature in this chapter. We also validated the primary points of discussion with subject-matter experts on our research team, the broader AIDE Initiative team, and the remarks of emergency managers who participated in a series of interviews, discussions, and workshops conducted by Aspen Digital as part of the AIDE Initiative. See Appendix B for additional details.

[2] Fichman and Kemerer, 1997; Damanpour and Schneider, 2006; Gallivan, 2001; Rogers, 2003; Tolbert and Zucker, 1983; Zhu, Kraemer, and Xu, 2006.

[3] Argonne National Laboratory has published the most comprehensive and most current information available about SLTT EM (National Preparedness Analytics Center, 2025). However, *current* is used in a research sense; at the time of this writing, more than a year has passed since those data were collected and published. SLTT EM access to and use of AI may have increased since the report was published.

[4] National Preparedness Analytics Center, 2025.

eight technological resources surveyed for all levels of government represented in the data, and it is an in-house resource at the same low rate at each level (Figure 4.1).

Among local offices, the largest "cannot access" rates are concentrated among rural, suburban, county, and small population (fewer than 50,000 people) jurisdictions, all at 39 percent.[5] A 2025 report published by Deloitte finds that no state can be described as "cutting edge" regarding technology; AI is the technology that state agencies most want but do not yet have.[6]

The barriers associated with moving from technology interest to implementation—whether AI or another technology—are significant. SLTT respondents identified a lack of funding and limited staff expertise or training as the two primary obstacles; the state-level barriers were affirmed in another study conducted within the same period.[7] Local offices also frequently cited limited knowledge of available resources and difficulty justifying return on investment.[8] State agencies frequently also cited privacy, security, and interop-

FIGURE 4.1
Local and State EM AI Access

Indicator	Local EM	State EM	What it suggests
In-house AI resource	11%	11%	Dedicated internal AI capacity remains uncommon
Any AI access (in house or other)	23%	58%	States report much higher overall access
No AI access and unable to obtain	37%	25%	Significant access barriers remain
Unsure whether office has AI access	22%	17%	Awareness and governance gaps persist
Want AI or ML but not currently in use	—	79%	Strong unmet demand at the state level

SOURCES: Features information from National Preparedness Analytics Center, 2025; Basu and Sheets, 2025.
NOTE: Percentages reflect the share of surveyed offices reporting each condition. State EM "Want AI/ML" reflects a distinct survey question asked in Basu and Sheets, 2025; local offices were not asked this question.

[5] National Preparedness Analytics Center, 2025.

[6] Ragini Basu and Trina Sheets, *Deloitte–NEMA National Risk Study 2025: Changing Landscapes in State Emergency Management*, Deloitte Center for Government Insights, Deloitte, 2025.

[7] National Preparedness Analytics Center, 2025; Basu and Sheets, 2025.

[8] National Preparedness Analytics Center, 2025.

erability concerns.[9] These findings are notable not only because they are substantial but because they echo barriers documented in earlier decades of EM research.[10] Although technologies have changed, many of the organizational constraints affecting adoption appear to be persistent.

These findings help explain why many offices may struggle to move to implementation. In Chapter 3, we documented that most products require a level of technical integration that presupposes IT capacity. Many products require ongoing data input. Fewer than one-third of technology companies offer 24/7 support. Pricing is opaque enough to make the full cost of implementation hard to know in advance. Most products carry some level of risk. Even offices that reach implementation may struggle to sustain product use because of turnover and skill gaps, making the achievement of standardization difficult.[11]

Not all EM offices will face these constraints equally. Larger population jurisdictions reported markedly higher AI access rates than smaller ones, and local offices serving populations with more than 500,000 residents were more than twice as likely to report access to AI resources compared with those serving fewer than 50,000 residents.[12] Many smaller offices also operate with one or fewer full-time staff and rely on part-time, volunteer, or dual-hatted personnel, and state offices report their own staffing shortages and funding pressures.[13]

All things considered, getting to implementation and from there to full adoption is likely to proceed faster in EM offices with stronger staffing, sufficient levels of IT and technical skills, and funding support. Offices operating under long-standing constraints will likely face a steeper path unless ways to overcome these challenges are identified and undertaken.

Diffusion Is Possible in EM, but the Conditions That Enable It Are Not Available Yet for AI

In Chapter 1, we distinguished the intraorganizational adoption process (Figure 1.2) from the field-level diffusion process (Figure 1.3). The previous section in this chapter dealt with the first process; this section deals with the second process. Diffusion in EM offices is not unprecedented. For EM-specific products, some products have diffused widely across local and state EM offices despite all the documented barriers (e.g., social media platforms, public alert and warning systems, virtual EOCs, and GIS).[14] Although prior research does not tell us how robustly these products are used or to what end, there is evidence that the broad diffusion of technology in EM is possible and has precedent.

---

[9] National Preparedness Analytics Center, 2025.

[10] See, for example, Thomas E. Drabek, "Microcomputer Usage in Disaster Preparedness and Response," *Industrial Crisis Quarterly*, Vol. 5, No. 2, June 1991; Eliot A. Jennings, Sudha Arlikatti, Simon A. Andrew, and KyungWoo Kim, "Adoption of Information and Communication Technologies (ICTs) by Local Emergency Management Agencies in the United States," *International Review of Public Administration*, Vol. 22, No. 2, 2017; Nicolas LaLone, Phoebe O. Toups Dugas, and Bryan Semaan, "The Technology Crisis in US-Based Emergency Management: Toward a Well-Connected Future," in Tung X. Bui, ed., *Proceedings of the 56th Hawaii International Conference on System Sciences*, University of Hawaii, 2023; Christopher Reddick, "Information Technology and Emergency Management: Preparedness and Planning in US States," *Disasters*, Vol. 35, No. 1, January 2011.

[11] For example, National Preparedness Analytics Center, 2025; LaLone, Toups Dugas, and Semaan, 2023.

[12] National Preparedness Analytics Center, 2025.

[13] National Preparedness Analytics Center, 2025; Basu and Sheets, 2025.

[14] National Preparedness Analytics Center, 2025; Basu and Sheets, 2025.

Early adoption is the first mechanism in Figure 1.3, and survey evidence from 2025 suggests that there is enough adoption across local and state EM offices to potentially generate observable evidence for others.[15] As we noted above, however, adoption is concentrated among larger and better-resourced offices. Without change, diffusion is likely to occur primarily within those higher-capacity offices.

The infrastructure that could support demonstration and uncertainty reduction—additional diffusion mechanisms—is limited. Although the National Research Council recommended "honest broker" technology clearinghouses (what it called "Consumer Reports for disaster managers") in 2007, no field-wide clearinghouse for AI-enabled EM products has emerged.[16] FEMA previously made products available to SLTTs and others who support EM functions.[17] It is unclear whether the agency will support AI product development going forward.

For peer effects to work, there must be functioning networks through which learning spreads. The International Association of Emergency Managers, Big City Emergency Managers, and the National Emergency Management Association, in addition to state-level EM professional organizations, may be vehicles for such dissemination (for example, through the Emerging Technologies Caucus–hosted lunch and learn series and EMPowerment e-Learning Series).

Institutional pressures shape adoption when organizations align with practices that are perceived as legitimate, appropriate, or expected within their field. Norms, standards, and regulatory expectations emerge over time and reinforce adoption. For AI in EM, these pressures and norms do not yet appear fully developed. SLTT AI governance and policy are in development now, and how both manifest across the country will either facilitate or hinder AI adoption.[18] A key condition that enabled diffusion for earlier technologies, chief among them active state-to-local provisioning, is not yet mature and positioned for AI in the near term.[19] Only 3 percent of state EM offices make AI resources available to local or tribal EM offices. This is the lowest provisioning rate for any technology surveyed.[20] The state-as-backstop story is also weaker than it appears. State EM offices face severe workforce strain, including turnover and insufficient staffing and funding.[21]

It is possible that institutional pressures related to AI-enabled products could emerge from the parent organizations in which many SLTT EM offices sit. Parent organizations could procure AI-enabled products and expect that suborganizations, such as EM offices, use those products. Otherwise, past EM technology diffusion patterns suggest that the diffusion of AI-enabled products is most likely to occur for the most basic and broadly useful applications: products that are on par with computers and email. Research on technology diffusion shows that basic and broadly useful products tend to diffuse more widely, while sophisticated

---

[15] National Preparedness Analytics Center, 2025.

[16] National Research Council, *Improving Disaster Management: The Role of IT in Mitigation, Preparedness, Response, and Recovery*, National Academies Press, 2007, p. 76.

[17] FEMA, "Resilience Analysis and Planning Tool (RAPT)," webpage, last updated April 8, 2026; FEMA, "FEMA Geospatial Resource Center," webpage, undated.

[18] Sha Sajadieh, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Lapo Santarlasci, Juan Pava, Nestor Maslej, Russ Altman, Erik Brynjolfsson, et al., *Artificial Intelligence Index Report 2026*, AI Index Steering Committee, Institute for Human-Centered Artificial Intelligence, Stanford University, April 2026; Amy Glasscock, Emily Lane, Eric Sweden, Alex Whitaker, and Kalea Young-Gibson, *The 2025 State CIO Survey: Leading Change Through Uncertain Times*, National Association of State Chief Information Officers, 2025.

[19] State EM offices have historically helped some local governments access selected technologies and shared systems, although available evidence suggests meaningful variation across states in the specific technologies provided.

[20] National Preparedness Analytics Center, 2025.

[21] National Preparedness Analytics Center, 2025; Basu and Sheets, 2025.

products often cluster in better-resourced agencies.[22] Given the current conditions for local and state EM, the broad diffusion of AI-enabled products across these offices is possible, but it is unlikely to occur soon or evenly without some novel interventions.

AI Adoption Across Organizations That Support EM Is Meaningful but Emerging at Different Rates

Many products in the landscape would likely be used primarily by organizations other than EM offices, including utilities, nonprofits, public works departments, public health agencies, transportation organizations, schools, hospitals, law enforcement, fire services, and communications providers, which support EM task areas across preparedness, response, and recovery. Whether AI-enabled products become broadly used across those fields is therefore also relevant to the future of EM practice and associated outcomes, even if such use is analytically distinct from direct EM-office adoption.

The same Chapter 1 adoption frameworks apply outside EM offices. Diffusion would still hinge on individual organizations fully adopting AI products. We would still be looking for the same benchmarks as signs of field diffusion. The available evidence suggests that adoption in both government and nonprofit settings is meaningful and expanding but still uneven across sectors, organizations, and functions.[23]

Across SLTT governments, several surveys suggest meaningful workplace use of AI, and state chief information officer organizations reported more than 50 percent year-over-year growth in generative AI use.[24] Use varies considerably by function and occupation: IT professionals reported higher use than other SLTT government professionals, and reported priority use cases—administrative support, public communication, constituent engagement, cybersecurity and fraud detection, and selected analytics—differed across states, counties, and cities.[25] The pattern suggests progress but not a uniform whole-of-government transformation.

The nonprofit sector shows a similar pattern.[26] Larger nonprofits with greater revenue and staffing are overrepresented among the early adopters,[27] while smaller organizations more commonly cite budget, capability, and training constraints.[28]

Taken together, the available evidence suggests that some organizations adjacent to EM appear further along in AI experimentation or selected implementation than EM offices, especially for common lower-friction uses, such as drafting, communication support, document processing, customer or constituent

---

[22] For example, Rogers, 2003.

[23] Center for Digital Government, *2024 Generative AI Survey*, e.Republic, May 2024; Glasscock et al., 2025; Gerald Young, *Artificial Intelligence in the Workforce: A Survey of State and Local Employees*, MissionSquare Research Institute, June 2025; Janne Rotter and William Bailkoski, "AI Adoption in NGOs: A Systematic Literature Review," arXiv, arXiv:2510.15509v1, October 17, 2025; ServiceNow, *Impact AI: Nonprofit Digital Transformation Report 2024*, 2024; Wanzhu Shi and Lauren Azevedo, "Determinants of AI Adoption in Nonprofit Organizations," *Nonprofit and Voluntary Sector Quarterly*, 2026; Virtuous and Fundraising.AI, *The 2026 Nonprofit AI Adoption Report: A Benchmark Study of 346 Organizations*, Virtuous, February 2026.

[24] Center for Digital Government, 2024; Glasscock et al., 2025; Young, 2025.

[25] Center for Digital Government, 2024; Young, 2025.

[26] Sudheekar Reddy Pothireddy, "From Data to Impact: How AI and Microsoft Enhance Efficiency in Nonprofit Organizations," *International Journal for Multidisciplinary Research*, Vol. 7, No. 1, January–February 2025; Rotter and Bailkoski, 2025.

[27] ServiceNow, 2024; Rotter and Bailkoski, 2025.

[28] Virtuous and Fundraising.AI, 2026.

engagement, and workflow efficiency.[29] That relative position likely reflects structural differences rather than sectoral enthusiasm alone. Many adjacent organizations operate with larger budgets, dedicated IT functions, stronger enterprise software environments, clearer return-on-investment cases, or more-continuous daily operations than many EM offices.[30] This should not be interpreted to mean that these sectors have solved adoption challenges or moved broadly into advanced use.[31]

The documented barriers for EM offices recur across government and nonprofit AI adoption (e.g., funding, staff capacity, data and governance concerns, and procurement frictions), although the relative severity of these barriers varies by sector, organization, and location.[32]

At the field level, however, these barriers do not appear equally binding across all settings because the mechanisms that support diffusion are stronger in some environments than others. SLTT governments appear to possess several documented diffusion mechanisms, including peer-to-peer learning, professional networks, cooperative purchasing authorities, technology company and consultant engagement, visible early adopters, and cross-jurisdictional initiatives, but, again, the presence of these mechanisms varies by sector and subsector.[33] Prior research on chatbot adoption in government provides one example of how a once-limited AI capability spread, aided in part by operational pressures during the coronavirus disease 2019 pandemic.[34] Diffusion mechanisms are also present in the nonprofit sector (e.g., early adopters or pacesetters and various peer networks).[35] Other hallmarks of diffusion are not yet present in this sector.[36]

Both government organizations and nonprofits may benefit from the relatively low entry cost and accessibility of many generative AI products, as well as from growing field attention to AI, but the evidence also indicates persistent resource, funding, training, and governance gaps that may limit broader institutionalization.[37]

---

[29] Center for Digital Government, 2024; Glasscock et al., 2025; Young, 2025; Virtuous and Fundraising.AI, 2026; Shi and Azevedo, 2026; Elisha Smith Arrillaga, Seara Grundhoefer, and Christina Im, *AI with Purpose: How Foundations and Nonprofits Are Thinking About and Using Artificial Intelligence*, Center for Effective Philanthropy, 2025; Sarah Di Troia, Vanessa Parli, Juan N. Pava, Haifa Badi Uz Zaman, and Kelly Fitzsimmons, "Inspiring Action: Identifying the Social Sector AI Opportunity Gap," Institute for Human-Centered Artificial Intelligence, Stanford University, and Project Evident, Working Paper, February 2024.

[30] National Preparedness Analytics Center, 2025; Sara Hinkley, *Technology in the Public Sector and the Future of Government Work*, UC Berkeley Labor Center, University of California, Berkeley, January 2023; ServiceNow, 2024; Rotter and Bailkoski, 2025; Virtuous and Fundraising.AI, 2026; Ishana Ratan, Alison E. Post, Tanu Kumar, and Mridang Sheth, "Which Local Governments Adopt New Technology First? Agency Size and Organizational Champions for Open Transit Data," *Urban Affairs Review*, Vol. 62, No. 62, March 2026.

[31] Center for Digital Government, 2024; Glasscock et al., 2025; Virtuous and Fundraising.AI, 2026; Rotter and Bailkoski, 2025.

[32] National Preparedness Analytics Center, 2025; Center for Digital Government, 2024; Glasscock et al., 2025; Rotter and Bailkoski, 2025; Virtuous and Fundraising.AI, 2026.

[33] Center for Digital Government, 2024; Hinkley, 2023; Nari Johnson, Elise Silva, Harrison Leon, Motahhare Eslami, Beth Schwanke, Ravit Dotan, and Hoda Heidari, "Legacy Procurement Practices Shape How U.S. Cities Govern AI: Understanding Government Employees' Practices, Challenges, and Needs," *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, Association for Computing Machinery, 2025; Glasscock et al., 2025; National Association of State Procurement Officials, *2024 Survey of State Procurement Practices Report*, 2024; Ratan et al., 2026.

[34] See Tzuhao Chen, Mila Gascó-Hernandez, and Marc Esteve, "The Adoption and Implementation of Artificial Intelligence Chatbots in Public Organizations: Evidence from U.S. State Governments," *American Review of Public Administration*, Vol. 54, No. 3, April 2024.

[35] See ServiceNow, 2024.

[36] See Di Troia et al., 2024

[37] Steven Kwamena Aikins and Tamara Dimitrijevska-Markoski, "Determinants of Government AI Adoption for e-Government," *AI & Society*, 2026; Smith Arrillaga, Grundhoefer, and Im, 2025; Anne David, Tan Yigitcanlar, Rita Yi Man Li, Juan M. Corchado, Pauline Hope Cheong, Karen Mossberger, and Rashid Mehmood, "Understanding Local Government Digital Technology Adoption Strategies: A PRISMA Review," *Sustainability*, Vol. 15, No. 12, 2023; Lu Gao and Luis Felipe

These patterns carry two implications for EM. First, AI capabilities that are relevant to functional and task area outcomes may continue to enter the broader EM system through adjacent organizations, even if direct EM office adoption remains low. Utilities, hospitals, transportation agencies, nonprofits, and public safety organizations may improve aspects of their own operations in ways that indirectly strengthen preparedness, response, or recovery when they leverage those products in the context of EM. Second, unevenness is likely to remain a defining feature of adoption. Just as EM offices vary in capacity, so too do the organizations around them. Some jurisdictions may benefit from a broader ecosystem of relatively capable partners, while others may face weaker surrounding capacity across multiple sectors at once.[38]

For EM specifically, these implications reinforce a central theme of this chapter: AI adoption is unlikely to emerge as a single, rapid, and uniform shift across all relevant organizations. AI adoption is more likely to proceed unevenly across organizations, often concentrated first in common administrative and communication uses and shaped heavily by institutional capacity, professional networks, available support, surrounding governance and policy, and operational need.

Technology Company Design Choices and Communication Practices Shape What Practitioners Can Do

In the previous sections, we examined adoption primarily through the lens of the organizations doing EM work. In this section, we shift the lens. Technology companies are not the entities adopting AI, but they materially influence the conditions under which others decide whether and how to adopt. Their product design choices, business model choices, support structures, and communication practices can influence how easy or difficult it is for potential buyers to move from awareness to standardization.

In Chapter 3, we emphasized that a product listed on the market is not the same as a product in use. The gap between availability and adoption is shaped by the structural features of the product, the technical and operational work required to stand up the product and keep it running, what the product costs, and what risks the product carries. Public marketing materials do not always make these burdens visible, leaving buyers to identify them during evaluation or procurement.

Many products implicitly assume that a customer has the procurement bandwidth, implementation funding, and ongoing data-provision capacity to stand them up and keep them running. Where those assumptions do not hold, adoption may be out of reach regardless of practitioner interest. The barrier is not always price: A product that is offered at no cost but requires a significant level of technical capacity to deploy may be less accessible in practice than a moderately priced product that arrives ready to use.

There is an opportunity for technology companies to make products easier to find and evaluate for the EM buyer who is short on time and resources.

Technology company communication practices shape adoption upstream of procurement. There is an opportu-

Luna-Reyes, "Implications of Adopting AI-Based Applications in Local Governments: A Systematic Literature Review from a Public Value Perspective," in Younhee Kim and Michael J. Ahn, eds., *The Art of Digital Governance: Navigating Platforms and AI Revolution*, Springer, 2025; Di Troia et al., 2024; ServiceNow, 2024; Virtuous and Fundraising.AI, 2026; Rotter and Bailkoski, 2025.

[38] National Preparedness Analytics Center, 2025; Center for Digital Government, 2024; ServiceNow, 2024; Virtuous and Fundraising.AI, 2026; Rotter and Bailkoski, 2025.

nity for technology companies to make products easier to find and evaluate for the EM buyer who is short on time and resources. Standardized product disclosures—analogous to nutrition labels on food or window stickers on cars—could provide the clarity practitioners need. Technology companies can provide transparency about their own products and work toward standardization across the AI-enabled product industry. Transparent communication from technology companies would help practitioners evaluate products and move from evaluation toward adoption decision, the next step toward full adoption. There is also opportunity for technology companies, particularly big tech and hyperscalers with significant resources, to adopt evidence-informed design criteria to create products intended for broad adoption and diffusion.[39] EM offices at the SLTT level contend with perennial challenges, chief among them inadequate funding, limited technical expertise among staff, and insufficient staffing levels. These challenges may deepen as changing funding policies at the federal level affect available resources within SLTT EM offices. Fields outside EM face similar challenges to varying degrees. A product designed for broad adoption and diffusion would meet emergency managers and others who support EM functions where they are at. According to our analysis, products that would be best positioned for broad adoption and diffusion are those that would meet both task-related and technical criteria (see the box).

Evidence-Informed Design Criteria to Support Broad Adoption and Diffusion in EM

The following lists detail task-related and technical criteria that would support the broad adoption and diffusion of AI-enabled products in EM.

Task-related criteria

- The product supports all functional areas.
- The product supports more than one task area.
- The product supports emergency manager work related to a task area and the work of others who support that area.
- The product provides functionality to support coordination across groups involved in the task area.
- The product supports tasks across the breadth of a task area.
- The product supports the integration, synthesis, and access of data related to the task area.

Technical criteria

- The product requires no or little training to use.
- The product requires light or no integration.
- The product has no hidden add-ons or extensions.
- The product has no requirement for ongoing data input that the user must provide.
- The product can function with intermittent or no user internet access (response products).
- The product has low or no privacy, legal, or cyber risks for users.
- The product is from a company with cybersecurity and AI governance certifications.
- The product is cheap or free.
- The product has no hidden ongoing costs.

[39] By *big tech*, we mean the largest and most influential technology companies, whose platforms, capital, data, and market reach give them outsize power over digital markets and the direction of innovation. These firms often operate across multiple sectors, such as cloud computing, advertising, consumer devices, software, and AI.

Products evidencing the task-related criteria would be generally useful, and products evidencing the technical criteria could be adopted in low-capacity and high-restriction entities.

The technology market has shown that products with many of these characteristics can be developed to support general tasks (e.g., Anthropic's Claude, OpenAI's ChatGPT, Microsoft's Copilot). It is possible that technology companies might purpose-build products for EM or configure general products for EM tasks, either of which could be helpful depending on the task area the product is going to support.

Technology companies do not determine adoption alone, but their choices are a lever that shapes AI adoption outcomes, and those choices extend beyond product design and pricing. Technology companies that absorb some of the implementation burden—through managed deployment, preconfigured integrations, or dedicated onboarding support—reduce the total cost of adoption in ways that price reductions alone cannot. Companies that lower the full cost of adoption may reach a broader market and build durable relationships with emergency managers and others who support EM functions. Clearer product information in nontechnical language coupled with products that are designed for broad adoption and diffusion are distinct contributions that technology companies can make. Additional opportunities lie in more-integrated product offerings, clearer pathways for combining partial solutions, and reductions in the stacking burden across organizational EM workflows. Technology company behavior is therefore part of whether products become attainable, particularly for smaller and capacity-constrained organizations with the least margin for costly trial and error.

Institutional Conditions Shape What Practitioners Can Realistically Do

Practitioners operate within institutional systems whose rules, capacities, and decisions shape whether AI products can be evaluated, procured, implemented, and sustained. Four parts of that wider setting are especially relevant here: IT departments, procurement policies and processes, AI governance, and funding and staffing decisions. Many of the forces that shape these four parts of the environment are, like product information and design, outside the direct control of EM offices or the other organizations doing EM work.

Information Technology Capacity Often Sits Elsewhere and Is Stretched

In Chapter 3, we showed that most products harbor integration and ongoing data demands that presuppose certain levels of IT and technical capacity. For many EM offices, those capacities are not in-house. To the extent that those capacities exist, they sit in the surrounding government's IT department, where EM is one of a group of other internal customers competing for attention. This pattern has been documented across multiple generations of EM technology research.[40]

IT capacity is constrained at the state and local level. State chief information officers report among their leading challenges legacy systems, workforce shortages, funding pressures, procurement complexity, and expectation management regarding newer technologies, including AI.[41] Cybersecurity dominates many priorities, and the services that state chief information officer organizations most commonly extend to local governments are weighted toward security rather than support for new technology adoption.[42] Where AI-specific review structures exist, they vary considerably: Some cities have created dedicated approval pro-

[40] See Drabek, 1991; Jennings et al., 2017; LaLone, Toups Dugas, and Semaan, 2023.

[41] Glasscock et al., 2025.

[42] Glasscock et al., 2025.

cesses, others offer optional advisory support, and many appear to have neither.[43] For AI-enabled products to support the work of EM, whether enough enabling policy and IT capacity will be available, soon enough, and oriented to support AI-enabled product adoption is an open question. *Leaders of organizations and those who are responsible for IT will ultimately determine the extent to which AI-enabled product adoption and diffusion will occur and how quickly, not necessarily emergency managers or others who support EM functions. Therefore, advancement and solution pathways must include an assessment of the influence of leadership outside EM.*

Procurement Systems Often Lag AI Acquisition Realities

In Chapter 3, we documented complexities of procurement. Product pricing is often opaque, and even substantial research effort failed to yield basic cost information that we could have high confidence in for most products. The institutional patterns behind that complexity matter for adoption.

Procurement authority over technology is structured differently across SLTTs, affecting which products are attainable in each jurisdiction. Formal and small-purchase thresholds can vary substantially.[44] A purchase that requires competitive solicitation in one state may be acquired with a purchasing card in another.[45] AI-specific procurement infrastructure also appears limited. State adoption of AI-specific procurement terms sat at roughly 41 percent in 2025.[46]

At the local level, some AI acquisition may occur outside formal procurement pathways through purchasing cards, free subscriptions, donations, or piggyback contracts that rely on technology company terms negotiated elsewhere.[47] Local thresholds and oversight rules also vary.[48] As a result, the same AI product may face extensive review in one jurisdiction and limited review in another.[49] Even when an EM office identifies a promising product, the path to implementation may therefore be uneven and shaped by procurement systems that the office does not control.

Funding and Staffing Remain the Most Persistent Constraints

Research on EM offices, SLTT governments, and nonprofits consistently cites funding and staffing among the top barriers to AI product adoption. Although AI-enabled products are entering the market in increasing numbers, funding and staffing constraints remain critical barriers to adoption and diffusion.

These conditions can compound one another. Offices may lack the personnel to pursue funding opportunities and otherwise secure funding. Administrative and compliance work can consume scarce staff time, reducing what remains for mission activities. These constraints therefore affect not just getting to the point of full adoption and diffusion but whether organizations even move beyond the first step of awareness toward adoption. Leaders of organizations and local and state governments will decide whether adequate investment is made in EM offices and other organizations that support the work of EM.

---

[43] Johnson et al., 2025.

[44] National Association of State Procurement Officials, 2024; Jerrell D. Coggburn, "Exploring Differences in the American States' Procurement Practices," *Journal of Public Procurement*, Vol. 3, No. 1, 2003.

[45] National Association of State Procurement Officials, 2024.

[46] Glasscock et al., 2025.

[47] Johnson et al., 2025.

[48] Johnson et al., 2025; National Association of State Procurement Officials, 2024.

[49] Johnson et al., 2025.

The AI Policy Environment Is Still Forming

Funding decisions and broader AI policy direction flow through legislatures, executives, and the systems they create. The federal government, SLTTs, and nonprofits are actively determining how they will manage access to and the use of AI products and the infrastructure on which many such products depend (e.g., data centers). The assurance ecosystem for cybersecurity is still developing and only just emerging for AI. In Chapter 3, we documented that fewer than 1 percent of products in the landscape carry a technology company claim to either ISO/IEC 42001 or alignment with NIST's AI Risk Management Framework. By comparison, cybersecurity markets benefit from more-established signals, such as SOC 2, ISO 27001, and FedRAMP, although the requirements to attain these certifications are not universal. The decisions that are made about AI products by elected officials and the leaders of organizations will also have a determining influence on the extent to which both individual organization adoption and field diffusion occur within and across fields.

Conclusion

The central finding of this chapter is that EM-relevant AI adoption and diffusion are likely to proceed at varied paces across two pathways: (1) direct adoption within EM offices and (2) adoption by the many other organizations whose work supports EM outcomes, including utilities, hospitals, transportation agencies, and nonprofits, followed by diffusion within those fields. Progress will be shaped by not only whether useful products exist (they do, and more are coming) but also the choices and investments that people make. How quickly and broadly adoption and diffusion proceed and in what fields will depend on (1) technology companies' choices about what information they provide and how they design products and (2) policymaker decisions. When the decisions of these entities support access to and the use of AI-enabled products, progress is more likely to move beyond awareness to full adoption, and diffusion may follow. There are other initiatives that professional associations, philanthropy, big tech, and potentially FEMA can pursue to facilitate AI-product adoption and diffusion. At the same time, not all progress depends on forces beyond the personnel who do the work of EM, whether in EM offices or external to those offices. Emergency managers and others who support EM have levers they can use now to ready themselves to access and use AI when the conditions allow.

Broad diffusion in any field is unlikely to result from any single actor or one-time intervention. Diffusion is more likely to emerge through cumulative actions by multiple players who make AI-enabled products more practical, accessible, and useful over time, improving disaster preparedness, response, and recovery outcomes. In Chapter 5, we turn from analysis to facilitating action by outlining practical next steps for each stakeholder group.

CHAPTER 5

From Awareness to Broad Adoption and from Adoption to Diffusion: An Agenda for Action

Disasters are getting more frequent, more complex, and more costly, while emergency managers and others who support EM functions continue to face persistent resource constraints. The challenges are not new. The federal government's recent policy to shift more of the financial and managerial burden of EM to SLTT governments compounds these challenges, especially in the near term. Against that backdrop, this report finds reason for measured optimism.

AI-enabled products will not resolve the challenges emergency managers and others who support EM functions face, but the evidence in Chapter 2 indicates that there are products available with the potential to help EM personnel work faster, at greater scale, and with better information. The capabilities that emergency managers need most, such as situational assessment, faster public information and warning delivery, administrative relief, and training scenarios, are reflected in real products that are available now.

However, in Chapters 3 and 4, we documented that product availability alone does not ensure meaningful adoption. In those chapters, we identified barriers on the product side and the user side that may slow broad adoption and diffusion of AI-enabled products. Additionally, organizational and cultural change factors reinforce the notion that procurement alone will not produce adoption. Government agencies, nonprofits, utilities, and other EM-supporting organizations may be risk-averse and slow to adopt new technologies, especially when those technologies affect public safety, legal exposure, resource allocation, or public trust. As we identified earlier, other adoption considerations could include leadership buy-in, workflow integration, human-oversight norms, user training, and transition planning from product pilots to sustained use. Yet, although there are barriers, many of those barriers are solvable. Broader gains are achievable if key stakeholders improve the surrounding conditions together.

A Sequenced Agenda for Action

Improving the conditions for AI adoption and diffusion requires system change rather than the work of any single actor. EM is, at its core, a distributed function; so, too, is the adoption and diffusion of AI within it. Technology companies can make products easier to evaluate, procure, and integrate. Policymakers can address the IT, procurement, governance, and funding conditions that determine what reaches the field. Researchers can answer open questions about how AI products perform against the criteria that matter for hazards and disaster work. The work can begin now.

We provide each of these key stakeholder groups with a sequenced agenda for action (Appendix G) that identifies practical things each group can do to support progress over three time horizons. The first horizon is **what should happen now** (the next 12 months): practical steps that can reduce friction, improve readiness, and accelerate near-term progress. The second horizon is **what should be built in the near term** (the next 1–3 years): harder but important efforts to strengthen enabling systems, such as procurement pathways, governance, and workforce capability. The third horizon is **what should be** done on an ongoing basis: generat-

Many of the barriers to adoption and diffusion are not purely technical; they are problems of funding, knowledge and skill, accessibility, and policy

ing evidence to inform policy and practice, documenting what works, identifying recurring barriers, and guiding continuous improvement.

We do not claim that any single action will transform the adoption and diffusion of AI-enabled products that are relevant to EM. Nor do we suggest that, even if all the recommendations are pursued, adoption will become universal across EM offices or that diffusion will occur across all fields that support EM functions. We do conclude, however, that these stakeholder groups will shape the breadth and speed with which AI-enabled products come to support EM across preparedness, response, and recovery. The influence of these groups will be felt under the status quo and in scenarios in which they take deliberate action to accelerate adoption and diffusion.

Beyond these central stakeholder groups, philanthropic organizations, professional associations, and other intermediaries are uniquely positioned to support progress across all three time horizons and stakeholder groups. These entities can operate across organizational boundaries, align incentives, and invest in activities that other stakeholders are not structured or resourced to undertake. *The findings in Chapters 2 through 4 suggest that many of the barriers to adoption and diffusion are not purely technical; they are problems of funding, knowledge and skill, accessibility, and policy.* Philanthropic organizations, professional associations, and other intermediaries can drive stakeholder focus and commitment to issues of adoption and diffusion and invest in the action agenda items outlined in Appendix G in ways that are consistent with their missions, priorities, and resources. It will take time for philanthropic organizations, professional associations, and other intermediaries to determine which specific agenda items they want to support and for which stakeholder groups, but they do not have to wait to drive progress. These entities can also catalyze high-impact initiatives to accelerate adoption and diffusion at scale now.

High-Impact Initiatives to Accelerate Adoption and Diffusion at Scale

Technology companies, including big tech and others, and philanthropic organizations have an opportunity to make a difference now by jointly investing in a small set of high-impact initiatives to accelerate adoption and diffusion at scale in the near term. *To achieve meaningful impact, these efforts should be implemented in close collaboration with emergency managers, the broader set of organizations that perform EM functions, and policymakers.* At least four such initiatives are feasible within the next 12 months.

Launch an EM AI Readiness and Adoption Accelerator

Establish a national-level EM AI readiness and adoption accelerator to function as an objective, trusted Consumer Reports–style clearinghouse for AI-enabled EM products.[1] This effort would provide product information to the public, facilitating technology company–verification of product capabilities and inform-

[1] The U.S. Department of Homeland Security's System Assessment and Validation for Emergency Responders (SAVER) Program is a Consumer Reports–style organization that could serve as a model for the accelerator. SAVER conducts impartial, practitioner-relevant, and operationally oriented assessments and validations of emergency response equipment and provides federal and SLTT emergency responders with the information they need to make knowledgeable decisions about acquiring, using, and maintaining their equipment. SAVER reports focus on two questions: What equipment is available? and How does it perform? The reports can be found at Science and Technology Directorate, U.S. Department of Homeland Security, "System Assessment and Validation for Emergency Responders (SAVER)," webpage, last updated June 10, 2026.

ing buyers, thereby improving awareness among those who support EM and strengthening the evidence base that is available to policymakers. Such a clearinghouse should provide, among other things, independent product evaluations, standardized buyer guides, model procurement and contracting language, peer-learning support cohorts across jurisdictions, and practical deployment templates for low-capacity jurisdictions and organizations. The accelerator should also implement and evaluate product pilots, beginning with those that have low barriers to entry and align with the capacity constraints under which many emergency managers and EM-supporting organizations operate.

Support an Education and Training Program to Increase Awareness and Knowledge

To address the root barriers identified in the research—limited knowledge and technical skills—develop and sustain a national education and training program to increase awareness and technical capability among emergency managers and others who perform EM functions. Valuable topics the program could address include what AI is, how AI works, and key AI terminology; how to use LLMs, including what to do and what not to do; how to evaluate organizational readiness across data, workforce, governance, and infrastructure; how to evaluate AI-enabled products for potential procurement; how to develop vendor evaluation criteria, questions to ask before buying, and red flags to watch for; and implementation fundamentals, including standing up an AI product, managing the transition, and sustaining adoption. Make the program freely accessible; design the program as a tiered learning pathway for adult learners, offering foundational, intermediate, and advanced levels so that participants can build capability over time; and deliver the program in a way that recognizes the time and bandwidth constraints that are common to EM roles. Peer-to-peer learning through professional associations could be a complementary strategy to build awareness and knowledge; however, a formal, sustained program could be made available and marketed to all emergency managers and others who support EM functions without consideration of membership.

Spearhead the Development of a Standing Cross-Sector Alliance

Establish a standing cross-sector alliance to focus sustained attention on AI-enabled product adoption and diffusion and to advance solutions for the conditions that hinder adoption and diffusion. Ensure that anchor members with the capacity to financially support and sustain the alliance are in place from the outset—particularly big tech, hyperscalers, technology companies that are already building products for EM or products used in EM, and philanthropic organizations. The alliance can serve as a coordinating body to align incentives, share lessons, and maintain momentum across otherwise fragmented efforts. We do not suggest that the alliance should have any financial or administrative relationship to the other three initiatives we recommend.

Address Adoption-Relevant Research Gaps

Building on this study, we identified research questions whose answers would inform future actions to facilitate the adoption and diffusion of AI-enabled products. We specify these within the action agenda for researchers, including research questions that we think are particularly important to address: (1) why long-standing deployed products have not yet reached full adoption, (2) what outcomes result from the use of AI-enabled products in EM functions where such products have been implemented, and (3) whether national AI governance, procurement, and IT policies create systematic barriers to adoption and what policy changes could address those barriers. Producing answers to these questions will take time, but those answers will provide critical information to shape the work of the cross-sector alliance, inform potential pilot projects,

and provide evidence and potential solutions to policymakers. Investment now could expedite answers and therefore progress toward adoption and diffusion.

Conclusion

In this report, we documented that product availability does not equate to organizational adoption and adoption does not equate to field-wide diffusion. The distance between these groups is shaped by procurement complexity, opaque pricing, integration burdens, connectivity dependencies, governance gaps, and, most persistently, funding and staffing constraints for which no technology can substitute. Despite these realities, the market for AI-enabled EM products is still emerging. IT and procurement policies are not fixed, and AI governance is being developed. Policymakers can change policies or otherwise ensure that they allow under-resourced EM offices and those of the others who support EM functions to benefit from the promise of AI-enabled products. Interested technology companies could also communicate transparently to allow product evaluation without significant time investment and design their products with the lowest-capacity jurisdictions and organizations in mind. Philanthropic organizations, professional associations, and other intermediaries can hold the space, fund the work, and coordinate across boundaries that no single actor can cross. None of these actors can produce broad adoption and diffusion on their own, but each actor can move part of the effort forward. Together, by acting on the high-impact initiatives outlined in this chapter and the action agendas in Appendix G, they can make AI-enabled products a meaningful contribution to EM work, which is becoming increasingly demanding, and help overcome the accessibility barriers that smaller and low-capacity jurisdictions and organizations face.

The point of all of this is not to promote the use of AI-enabled products just because they exist. The point is that AI-enabled products—adopted and diffused—can help emergency managers and others who support EM functions save more lives, protect more property, and protect the environment. With these products, emergency managers and others who support EM functions can support the preservation of local culture, businesses, and economies when disasters strike and facilitate a whole and just recovery for affected communities. In a future where hazards and disasters are more frequent, more severe, more costly, and affect more communities, AI-enabled products can play a critical role in achieving these outcomes.

APPENDIX A

A Brief History of AI

A comprehensive history of AI is outside the scope of this report, but in this appendix, we provide a short summary intended to help readers make sense of the terminology, capabilities, and claims that they are likely to encounter when evaluating AI-enabled EM technologies. The goal is not technical depth but practical orientation—to provide enough context to read vendor materials, procurement documents, and the rest of this report with a clearer understanding of what AI is and is not.

Why AI Is Hard to Pin Down

One of the first challenges to understanding AI is that the term is a moving target. This is in part because there has never been consensus across philosophy, cognitive science, computer science, or even botany—on what *intelligence* means. A phenomenon called the *AI effect* captures this well: As soon as a computer reliably solves a problem that was once thought to require intelligence, the solution tends to be reclassified as just software.[1] Spam filtering, GPS routing, voice recognition, fraud detection, and autocomplete were all once considered cutting-edge AI; today they are unremarkable features of everyday products. The same may be true of many capabilities that are marketed as AI in the EM space.

For practitioners, the practical implication is straightforward: **The AI label says very little about what a system does**. A product described as AI-powered may rely on decades-old statistical methods, modern LLMs, or something in between. Effective evaluation requires looking past the label to understand the underlying capability, the data it depends on, and the tasks it is genuinely suited for.

How AI Has Evolved

AI has moved through several broad eras, each defined by a different approach to making machines perform tasks associated with human intelligence.[2] The earliest era, often called symbolic AI (roughly 1950–1980), encoded human knowledge as explicit rules and logical statements.[3] This is the era in which Alan Turing proposed the Turing test (still referenced today), John McCarthy coined the term *artificial intelligence*, and the first chatbot (ELIZA) was developed.[4] During this era, AI relied on logic systems, rule engines, search algo-

[1] Douglas R. Hofstadter, *Gödel, Escher, Bach: An Eternal Golden Braid*, Basic Books, 1979; Pamela McCorduck, *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*, 2nd ed., A K Peters, 2004.

[2] Russell and Norvig, 2022.

[3] A. M. Turing, *Proposals for Development in the Mathematics Division of an Automatic Computing Engine*, National Physical Laboratory, E.882, 1945.

[4] Russell and Norvig, 2022; John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," *AI Magazine*, Vol. 27, No. 4, [1955] 2006; Joseph

rithms, and theorem provers to complete reasoning tasks.[5] A related approach, expert systems, rose to prominence in the late 1970s and 1980s and attempted to replicate expert decisionmaking in narrow domains by capturing specialist knowledge as large sets of rules enabled by knowledge bases, inference engines, and forward and backward chaining.[6] These early approaches were limited by their dependence on human-authored rules and struggled with the messiness of real-world data.

In the 1990s and 2000s, the field shifted from telling computers *how* to think to letting them *learn* patterns from data.[7] This statistical and machine learning era introduced such techniques as decision trees, support vector machines, Bayesian models, and regression analysis that could make predictions from labeled datasets.[8] Many of the AI capabilities embedded in everyday software—spam filters, recommendation engines, fraud detection—emerged from this era.

The Eras Most Relevant to EM Today

The capabilities that EM practitioners are most likely to encounter in current vendor offerings come from three more-recent eras, which build on one another.

Deep Learning and Perception (2010s)

Advances in computing power, the availability of very large datasets, and improvements in a class of techniques called *neural networks* enabled major breakthroughs in tasks involving perception—interpreting images, audio, and language.[9] This is the era that produced practical CV, reliable speech recognition, and NLP capable of supporting voice assistants, such as Apple's Siri and Amazon's Alexa.[10] These capabilities were enabled by methods that include convolutional neural networks, recurrent neural networks, and backpropagation.[11] In the EM context, technologies that automatically classify imagery (e.g., damage assessment from drone or satellite photos), transcribe radio traffic, or detect anomalies in sensor data generally trace back to this era.

Weizenbaum, "ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine," *Communications of the ACM*, Vol. 9, No. 1, 1966.

[5] Nils Nilsson, *Artificial Intelligence: A New Synthesis*, Morgan Kaufmann, 1998.

[6] Edward A. Feigenbaum and Pamela McCorduck, *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, 1st ed., Addison-Wesley, 1983; Peter Jackson, *Introduction to Expert Systems*, 3rd ed., Addison-Wesley, 1998; Russell and Norvig, 2022, Chapter 7.1; Edward A. Feigenbaum, "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Vol. 2, 1977.

[7] Tom M. Mitchell, *Machine Learning*, 1st ed., McGraw-Hill, 1997; Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.

[8] Bishop, 2006.

[9] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, "Deep Learning," *Nature*, Vol. 521, 2015.

[10] LeCun, Bengio, and Hinton, 2015; Apple Machine Learning Research, "Hey Siri: An On-Device DNN-Powered Voice Trigger for Apple's Personal Assistant," October 1, 2017; Ruhi Sarikaya, "The Technology Behind Personal Digital Assistants: An Overview of the System Architecture and Key Components," *IEEE Signal Processing Magazine*, Vol. 34, No. 1, January 2017.

[11] LeCun, Bengio, and Hinton, 2015; David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams, "Learning Representations by Back-Propagating Errors," *Nature*, Vol. 323, 1986.

Generative AI (late 2010s to early 2020s)

Earlier AI primarily *recognized* or *classified* things, but generative AI *creates* new content (e.g., text, images, audio, video, code) that resembles what a human might produce.[12] This era is best known for LLMs, such as OpenAI's GPT-2 and GPT-3, which demonstrated that a single model trained on enormous amounts of text could generate coherent writing across a variety of tasks.[13] The underlying technical advances—an architecture called the *transformer* for language models and *diffusion models* for image generation—are what make modern generative AI possible.[14] In the EM context, generative AI is what powers the products that draft incident reports, summarize situational updates, generate public messaging, or answer natural-language questions about plans and procedures.

Foundation and Multimodal AI (early 2020s)

The most recent era is defined by very large *foundation models*—models trained on diverse data spanning text, images, audio, and sometimes video—that can be adapted to many different downstream tasks.[15] Such products as Claude, Gemini, GPT-4 (the underlying model), and ChatGPT (the consumer product) are built on foundation models.[16] The term *multimodal* refers to a model's ability to take in and produce more than one type of content within a single interaction—for example, accepting a photograph and a written question and returning a written answer. In the EM context, this is the technology behind the products that can ingest mixed inputs (a map, a situation report, and an emergency services transcript, for instance) and produce a unified summary or recommendation.

What This Means Going Forward

Viewed through the lens of the AI effect, many of the capabilities we described in this report are unlikely to broadly remain AI in the public imagination for long. Just as spam filtering and GPS routing are now treated as ordinary software features, the generative and foundation-model capabilities that are now entering EM technologies will likely become standard, embedded, and unremarkable within a relatively short period.

---

[12] Bommasani et al., 2021; Ian J. Goodfellow, Jean Pouget-Abadiey, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozairz, Aaron Courville, and Yoshua Bengio, "Generative Adversarial Nets," in Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger, eds., *Advances in Neural Information Processing Systems*, Vol. 27, 2014.

[13] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever, "Language Models Are Unsupervised Multitask Learners," OpenAI, 2019; Brown et al., 2020.

[14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin, "Attention Is All You Need," in Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, eds., *Advances in Neural Information Processing Systems*, Vol. 30, 2017; Jonathan Ho, Ajay Jain, and Pieter Abbeel, "Denoising Diffusion Probabilistic Models," in Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, *Advances in Neural Information Processing Systems*, Vol. 33, 2020; Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli, "Deep Unsupervised Learning Using Nonequilibrium Thermodynamics," in Francis Bach and David Blei, eds., *Proceedings of the 32nd International Conference on Machine Learning*, Vol. 37, 2015.

[15] Bommasani et al., 2021.

[16] Anthropic, "Introducing the Next Generation of Claude," webpage, March 4, 2024; Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, et al., "Gemini: A Family of Highly Capable Multimodal Models," arXiv, arXiv:2312.11805, December 19, 2023; Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al., "GPT-4 Technical Report," arXiv, arXiv:2303.08774, March 15, 2023; OpenAI, "Introducing ChatGPT," webpage, November 30, 2022.

If that occurs, the questions facing EM offices and others who support EM functions will shift too. Some of today's procurement-stage challenges—distinguishing genuine AI capability from marketing language, evaluating vendor claims, deciding whether to adopt AI at all—may recede as these capabilities become ubiquitous and assumed. In their place, implementation and AI safety questions are likely to grow in importance: how these systems are configured and governed in operational settings, how their outputs are validated under time pressure, how errors and biases are surfaced and corrected, and how accountability is assigned when AI-assisted decisions affect public safety.

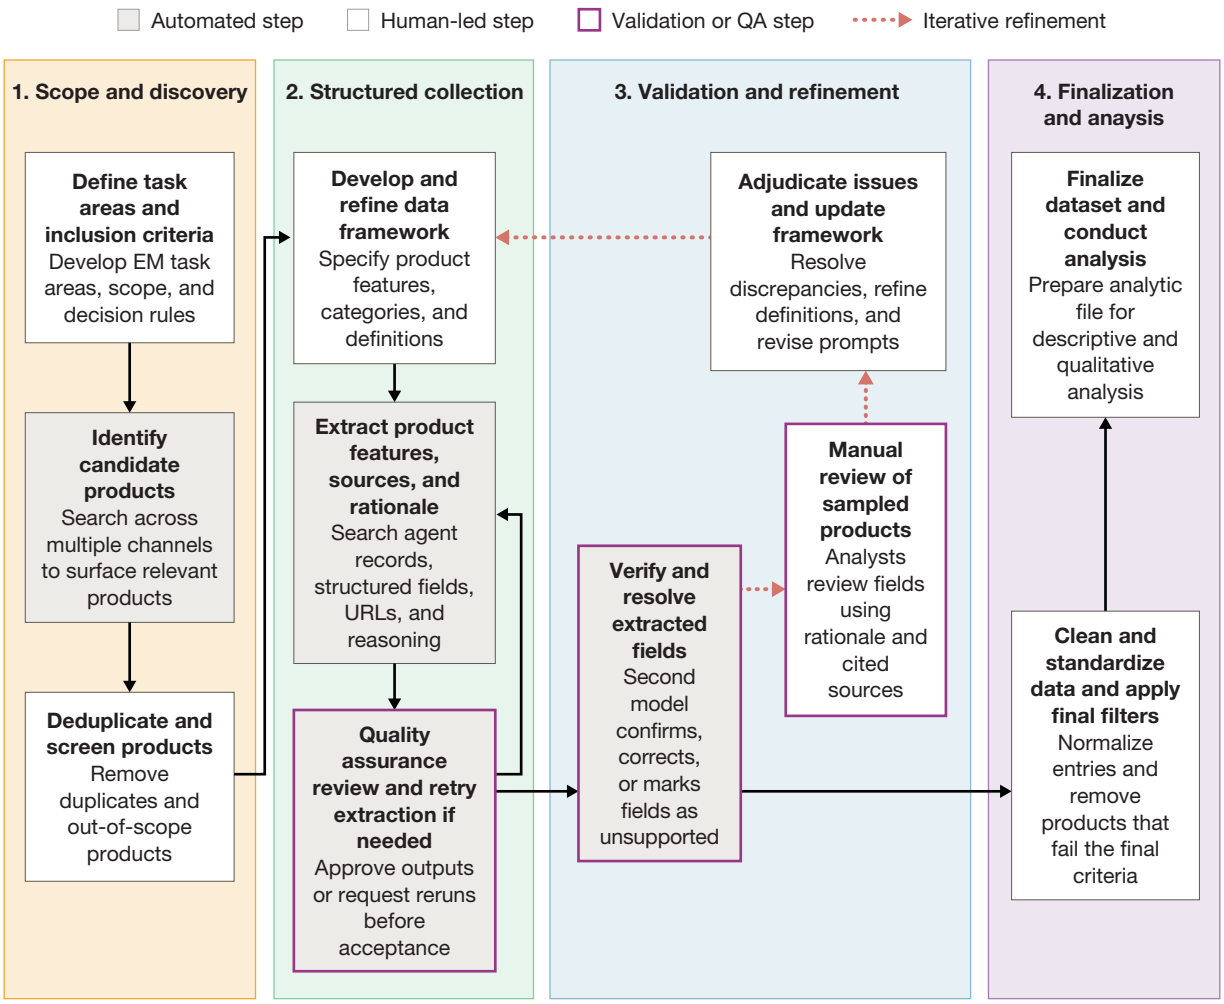
APPENDIX B

Project Approach

In this appendix, we describe the methods we used to produce the findings in this report and satisfy project deliverables. The work proceeded in five steps: developing the EM task areas that organized the research, identifying products on the market, developing a framework to guide information gathering, collecting and verifying product information, and analyzing the resulting dataset (Figure B.1). Each step is described in the following sections, followed by a statement on our use of AI in the research process and a discussion of limitations.

FIGURE B.1

Workflow for Product Identification, Information Extraction, Validation, and Analytic Dataset Preparation



Step 1: Develop EM Task Areas

Using authoritative EM documents—including the Emergency Management Accreditation Program Standards and the national planning frameworks—we identified 45 broad EM task areas: 20 response task areas, 16 recovery task areas, and nine preparedness task areas.[1] Task areas were validated through expert review and engagement with emergency managers.

Response task areas focus on operational functions, such as situational awareness, evacuation, emergency medical services, sheltering, logistics, and incident coordination. Recovery task areas include infrastructure restoration, housing, financial recovery, governance, and community engagement. Preparedness task areas include governance, risk assessment, planning, training, and exercises. The full list of task areas is provided in Appendixes C through E.

Step 2: Identify Products

We sought to answer three primary questions: (1) What AI-enabled or AI-related products are currently on the market to support EM use cases? (2) What potential exists for adoption and diffusion of these products to evolve EM practice? and (3) What must happen for this evolution to occur?

We searched for products across each of the 45 task areas using complementary strategies: (1) reviewing technology company lists from major EM and disaster conferences that took place from 2023 to 2025, (2) conducting two web searches, including iterative use of LLM APIs, (3) identifying relevant companies and products through a commercial company data API, and (4) obtaining stakeholder and expert recommendations.[2] The web searches were limited to publicly accessible sources. Company identification through the commercial API was conducted programmatically through the provider's documented search endpoint rather than through scraping or automation of any social media platform. The initial search surfaced 1,892 candidate products. After rounds of deduplication following each EM phase (response, recovery, preparedness) and the application of scoping filters (e.g., U.S.-headquartered technology company, contains AI, active product), the final dataset contained 1,179 distinct products from 717 unique technology companies.

We applied the following scoping criteria to frame the search:

- *AI-enabled* products are those for which AI provides or meaningfully contributes to the product's adaptive, predictive, generative, or reasoning capability. Products with no verified AI component were excluded.
- Products must be currently available and commercially acquirable, whether or not there is evidence of product deployment in a disaster context.
- Products may require human-in-the-loop approval as part of their operation.
- Products must be standard, named, and productized offerings with documented features, commercial terms, and repeatable deployment. One-off bespoke professional services builds were excluded.
- Products may be hardware, software, or integrated systems and may be delivered as SaaS, on-prem, or edge deployments.

[1] Emergency Management Accreditation Program, *Emergency Management Standard*, May 2022; FEMA, 2016; FEMA, 2019; NIST, 2023; NIST, 2024.

[2] Company and product searches were conducted using Enrich Layer, a commercial data provider that maintains an indexed company database that is accessible through a REST API. All queries were made through Enrich Layer's company search endpoint in accordance with the provider's terms of use (Enrich Layer, homepage, undated).

- Products may be purchased or licensed under any model, from open-source arrangements to proprietary closed-source licenses.
- Products may or may not support low-connectivity operations.

Step 3: Develop Framework to Guide Information Gathering

Through an iterative process that involved preliminary searches to assess what product information was publicly available, discussions with technology companies, and consultation with experts familiar with relevant public and privately held data sources, we developed a framework of product features that could be gathered consistently across the landscape. The framework organized information into six areas:

1. **Product and technology company overview** established product identity and a technology company footprint, capturing product descriptions, technology company information (including U.S. or international headquarters), and descriptions of the underlying AI in each product.
2. **Intended EM fit** captured whether the product was purpose-built for EM, had documented use in EM contexts, or was developed for other purposes but applicable to EM task areas. This area also captured primary and additional task areas that the product supports, the type of hazard the product applies to, and whether the product integrates with other existing EM products.
3. **Deployment and access model** captured how the product is delivered (e.g., SaaS, on premises, hardware) and how it is licensed or sold (e.g., subscription, enterprise license, per user).
4. **Burden and usability** captured the operational lift required to adopt the product, including data input requirements, level of integration effort required, and the training and technical support available from the technology company.
5. **Risk, compliance, and security posture** captured whether there are legal, privacy, or cyber concerns associated with the technology company or product and whether the technology company publicly lists evidence of cybersecurity certifications.
6. **Other implementation considerations** captured product maturity (from pilot through global deployment) and connectivity requirements (offline capable, intermittent tolerant, or online only).

We did not know it yet, but the framework was not finalized before data collection began. Definitions, categories, prompting strategies, and variables were refined continuously throughout the project as we learned what information could be reliably gathered, where initial operationalizations did not capture the distinctions that mattered, and where LLM interpretation errors or classification inconsistencies required correction. New framework components were added as we identified questions worth asking that had not been anticipated at the outset. We also ceased to collect information about some product dimensions because we could not gather reliable data. This iterative refinement continued through the final stages of analysis, and the framework described in this report reflects the cumulative result of that process rather than an initial design applied unchanged.

Step 4: Collect and Verify Product Information

We developed, refined, and implemented an information collection method aligned with the framework, paired with an iterative, multistage verification process that combined automated products, expert review, and a formal quality assurance review of our work and the code used to implement the automated components. The process is described in the following sections.

Multi-Agent Search Process

For each product, a search agent (OpenAI GPT-5.2 Pro API with web search capability) conducted targeted web searches to collect information for blocks of related product features.[3] For each feature, the agent recorded the information it found and the URLs of all the websites it visited and documented its reasoning for the assessment it provided in a rationale field. A quality assurance agent (Azure OpenAI GPT 5.2 API) then reviewed the search agent's output and either approved it or requested a retry. If none of the three retry attempts for a block were approved, the quality assurance agent selected the best response across the three attempts. Once all blocks for a product were approved, the process moved to the next product.

Automated Verification Process

After the search process, each product feature underwent automated verification using an LLM from a different provider chosen to reduce the risk of correlated model biases. The verification agent (Anthropic Claude Opus 4.6 API with web search capability) assessed the validity of each collected feature through its own independent web search, recording the URLs it visited and the rationale for its assessment.

The verification proceeded in two stages. In the first stage, the verification agent searched for independent evidence supporting the collected information. If the information was confirmed, the field was marked as verified. If the information was not confirmed, the agent searched again using the URLs retained from the original search process. If that second attempt confirmed the information, the field was marked as verified with sources. If the information still could not be confirmed, the process moved to the second stage.

In the second stage, the verification agent conducted a new search to retrieve corrected information. If corrected information was found and confirmed, the field was updated. If the corrected information could not be confirmed through a follow-up search, the field was marked as not supported. This branching process ensured that each feature had multiple opportunities for confirmation before being flagged as unverifiable.

Manual Review Process

After the automated verification process, a team of seven research analysts conducted a manual review of a randomly selected sample of 240 products, which was determined using the total number of products we had at the time (e.g., pre-exclusions), a sample size calculator configured for a 90 percent confidence level, and a ≤10 percent margin of error. Random sampling was selected over stratified approaches for three reasons: (1) the classification variables that could serve as stratification dimensions (e.g., task area assignments, product type) were outputs of the automated process subject to validation, making stratification circular; (2) the same data framework was applied uniformly to all products regardless of functional area or task area, meaning that the verification task did not vary structurally across categories; and (3) the scale of the product landscape made more resource-intensive sampling designs infeasible within project timelines. The resulting sample was nonetheless large enough relative to the dataset to naturally capture variation across product types, task areas, and discovery sources.

We developed a protocol to provide a consistent methodology across reviewers and communicated it to them through an initial training session and periodic meetings. Each reviewer received a subset of the sampled products and reviewed each product's features individually, guided by the rationale and source URLs collected during the automated process. We chose reviewers with substantial technical expertise in AI sys-

[3] A *search agent* is a software program that connects to an LLM to conduct internet searches and interprets the results autonomously. We wrote extensive code that provided each agent with detailed context about the project and the product being researched and specific instructions governing what to search for, what categories to apply, and how to document its reasoning.

tems; information about AI components was frequently not stated explicitly on technology company websites, and reviewers needed to draw on their own knowledge of AI architectures, model types, and deployment patterns to assess whether the automated classifications were reasonable given the available evidence.

Reviewers first read the rationale to understand why the model selected its answer. They then opened each cited URL to confirm the accuracy of the information. If the information could be confirmed at the cited sources, the reviewer marked the field as confirmed. If the information was vague or uncertain, the reviewer flagged it. If the information was clearly incorrect, the reviewer marked it as incorrect. Ultimately, each reviewer's determination for each field was guided by our standard of whether we could reasonably expect a human to reach the same conclusion as the LLM given the available information.

Percentage agreement was used as the reliability measure across reviewers. Agreement reached 93 percent across the review team, and for each product feature, at least 90 percent of field-level judgments were confirmed as consistent with the available evidence. After each review cycle, we met to share findings and identify improvements to the search and verification process, which was updated iteratively. The scope of this validation and its implications for interpreting dataset quality are discussed in the "Limitations" section.

Iterative Refinement

We concluded after the first manual review cycle that the search and verification process was effective, and subsequent review cycles focused on targeted improvements. The automated process was producing accurate results that were broadly consistent with human expert judgment, but our growing familiarity with the landscape revealed opportunities to capture information with greater precision and nuance than the original definitions allowed. Several product features were rerun through the collection and verification pipeline after we determined that the original category definitions could be sharpened. The types of refinements included adding new category options to accommodate product attributes not anticipated in the original design, allowing multiple category selections per product where the data required it, adding additional context for the LLM, and adjusting default classification logic for such fields as legal and privacy concerns to better reflect the evidence patterns observed in the data. Three new features were also introduced after the initial review: whether the product is a real, currently available product; product maturity; and interoperability with other EM products.

Data Cleaning

The completed dataset was cleaned to standardize entries across each field. Standardization tasks included normalizing formatting and capitalization, mapping free-text entries to controlled category values, and resolving inconsistencies across multisource entries to produce a dataset suitable for quantitative analysis.

Once the data were compiled and cleaned, we applied five filters that we determined were nonnegotiable for the sponsor. First, given that we used an LLM to gather the names of products, we removed any products from our list that were not currently available for procurement. Second, we removed any products that we did not verify to have at least a partial AI component. Next, we identified products that we could not verify as applicable to any EM functional area or task area after a manual review of the gathered evidence about the products. Fourth, we filtered out products that we learned were decommissioned or sunset because they are no longer available on the market. Finally, we removed any product associated with a technology company that is not headquartered in the United States.

Cost Estimation

Cost was among the most challenging dimensions to assemble, and we pursued it through multiple approaches over the course of the project. Early in the research, we conducted extensive manual searches of individual product websites to determine whether technology company–published pricing could be collected at scale. When that proved insufficient, we explored whether private-sector data sources—commercial databases, market intelligence services, and other proprietary data products—could be purchased or otherwise accessed to provide structured pricing information. None of these approaches yielded data that were consistent, comprehensive, or comparable enough to support systematic analysis across the landscape.

With those avenues exhausted, we spent time manually researching pricing across technology company websites, procurement databases, and public contract records to determine whether reliable cost data could be assembled through public sources at scale. This effort, which involved hundreds of hours of direct researcher engagement with technology company materials and pricing sources, revealed the depth of the challenge: Most technology companies do not publish pricing publicly, pricing structures vary substantially across products, and the information that is available is scattered across local government meeting documents, cooperative contract schedules, and marketplace listings that reflect specific negotiated arrangements rather than generalizable figures. This manual investigation also surfaced structural questions that shaped other parts of the research—including whether products function independently or depend on other products and how add-on pricing relates to base product costs—because the act of trying to determine what a product costs frequently revealed that the product itself was a component of a larger multiproduct procurement.

Drawing on what these prior efforts revealed, we designed an automated pipeline as a final attempt to estimate procurement costs for each product across three organizational sizes: small (one user or device), medium (20 users or devices), and large (150 users or devices). For each product, the pipeline sought to identify the primary pricing basis (per organization, per user, per device, or usage-based) and estimate low and high annual costs. Where AI features were priced as separate add-on modules rather than included in the base product, the pipeline collected pricing for those modules separately. Costs were then normalized to the three organizational sizes using stated assumptions. The pipeline also collected a confidence assessment for each estimate, distinguishing between estimates derived from publicly available technology company pricing (high confidence) and those derived from market-analogy comparisons with similar products (low confidence).

The pipeline pursued pricing through a systematic sequence of sources: technology company pricing pages, Amazon Web Services and Azure cloud marketplaces, federal procurement systems (e.g., GSA Advantage, USAspending, SAM.gov), cooperative contract price schedules, state procurement documents, and third-party software review sites. Most technology companies in this landscape do not publish pricing publicly, and most are not listed in federal procurement systems. Where pricing could be found, it was most often in local government meeting documents or cooperative contract schedules that reflect specific negotiated arrangements for specific jurisdictions.

Despite cumulative investment across all of these approaches—early manual website searches, exploration of private-sector data sources, sustained manual research of public procurement records, and a purpose-built automated pipeline—we found the resulting cost estimates to be insufficiently reliable for quantitative analysis, and no specific cost figures are referenced in this report. The information captured through this process was instead used to characterize the difficulty of understanding pricing for products in the landscape, including challenges resulting from variation in pricing structures and the difficulty of assembling comparable cost data across the market.

Step 5: Data Analysis

We analyzed the dataset using both descriptive statistics and qualitative analysis. We reviewed literature and consulted with subject-matter experts while data analysis was ongoing.

Descriptive Statistics

We compiled descriptive statistics to characterize the product dataset across each product feature. Univariate analysis produced counts and percentages for each category within each variable across various subsets of the data. For features that allow multiple values per product—such as deployment model and underlying AI—we analyzed the frequency with which categories co-occurred (for example, how often hardware overlapped with edge computing or how often NLP appeared alongside ML) and the distribution of the number of values per product. These analyses provided insight into how flexibly products can be adapted across deployment contexts and EM workflows.

We used multivariate analysis to compile cross-tabulations across groupings of complementary variables, such as privacy, legal, and cyber concerns. This gave us a count and percentage for how often categories from different variables co-occurred, providing insight into how product features correlated with one another. Within the multivariate statistics, we analyzed how each of the 45 task areas and the three higher-level EM functional areas corresponded to categories across all other features, identifying overlaps and gaps in the coverage the market provides. This included analysis of the most and least well-served task areas within each functional area and analysis of task area co-occurrence within individual products within each EM functional area and across all three functional areas.

Unknown values were included as a category in the univariate statistics but excluded from the multivariate cross-tabulations, where they did not add substantive value.

All descriptive analyses were conducted in R using the tidyverse package ecosystem and in Python using pandas.[4]

Qualitative Analysis Using Rationale Fields

Each feature in the product dataset has a corresponding rationale column that captures the information and reasoning that led to the category assignment. We used these rationale fields for three analytical purposes.

First, we used them to understand what the features and their categories actually represent in practice. Despite careful definitional work, the information available on product websites varies in clarity and specificity. Multiple reviewers noted during the manual review process the difficulty of classifying products with human judgment, which provided useful context for interpreting the automated classifications. Our prompting during the rationale analysis was conducted with extensive care to mitigate preemptively biasing the LLM outputs toward any conclusions. This process allowed us to identify patterns across the 1,179 products at a pace and scale that would not have been feasible through manual review within the project's constraints.

[4] The *tidyverse* is a collection of open-source R programming packages designed for data science, including tools for data manipulation, visualization, and analysis. It is maintained by Posit (formerly RStudio) and is one of the most widely used analytical toolkits in the R programming community (Hadley Wickham, Mara Averick, Jennifer Bryan, Winston Chang, Lucy D'Agostino McGowan, Romain François, Garrett Grolemund, Alex Hayes, Lionel Henry, Jim Hester, et al., "Welcome to the Tidyverse," *Journal of Open Source Software*, Vol. 4, No. 43, 2019). *Pandas* is an open-source Python programming library for data manipulation and analysis, providing tools for working with structured datasets, such as spreadsheets and databases. It is one of the most widely used data analysis libraries in the Python programming community (Wes McKinney, "Data Structures for Statistical Computing in Python," in Stéfan van der Walt and Jarrod Millman, eds., *Proceedings of the 9th Python in Science Conference*, 2010).

Second, we used the rationale fields to surface data quality issues and flag entries for recategorization. Through an iterative process, we identified cases for which the automated classification was inconsistent with the evidence described in the rationale. For example, some products were classified as having legal concerns, but the rationales indicated that the relevant lawsuit had been dismissed. On review, we made decisions to adjust the dataset where the evidence clearly warranted it.

Third, we used the rationale fields to characterize the products within consequential feature categories. For example, the qualitative analysis of deployment model rationales revealed that drones are the most common hardware product in the landscape. For security evidence, the analysis produced a compilation of the most common certifications. These characterizations informed the interpretive findings presented in the report.

Literature Review

We reviewed academic and gray literature concurrently with analysis of the product dataset. The goal of the literature review was to identify the extent to which technology adoption and diffusion has occurred in EM and why (including the adoption and diffusion of AI-enabled technology specifically) and the extent to which AI-enabled technology adoption and diffusion has occurred outside EM and why. We then analyzed the findings of this research in the context of the findings from the landscape of available AI-enabled products to support EM to identify the potential for those products to be broadly adopted and then diffused across EM offices and others who support EM functions.

Subject-Matter Experts

Throughout the analysis process, we discussed our findings with the EM subject-matter experts on our research team and those involved in the broader AIDE Initiative. We used these experts to validate the conclusions we reached from analyzing the landscape of available products against the findings from our literature review.

Additionally, the Markle Foundation contracted Aspen Digital, another collaborative organization supporting the AIDE Initiative, to facilitate a series of interviews, discussions, and workshops with emergency managers across the country. We attended all of these events. Emergency managers' descriptions of their use of technology and AI-enabled products and the barriers preventing the adoption of those products was consistent with what the literature described.

Transparency Regarding AI Use

This research made extensive and deliberate use of AI tools throughout the research process, from product identification to data collection and verification to analysis and writing.

LLMs were used in a variety of capacities. LLMs assisted in the product identification process, conducting structured web searches to surface AI-enabled products relevant to the 45 EM task areas; this was one of five complementary search strategies. LLMs also conducted structured web research that produced the product feature dataset, searching technology company websites, procurement documents, cloud marketplaces, and government records to extract information about each product against the predefined framework. A separate LLM from a different provider served as the initial verification layer, independently assessing the accuracy of the extracted information through its own targeted web searches.

LLMs also assisted in the qualitative analysis of rationale fields, identifying patterns, surfacing data quality issues, and characterizing product features across the dataset. Specifically, LLMs were used to surface

dominant patterns and identify illustrative examples within the rationale fields. They were not used to generate the quantitative findings reported in this study because LLMs are unreliable for counting and basic arithmetic. When an LLM highlighted a data point relevant to the analysis, we conducted independent verification using structured analytical techniques in R and Python. In these cases, LLMs were used to generate the analytical code, which was reviewed through RAND's quality assurance process and/or was tested to confirm it produced the expected results.

Finally, we used LLMs during the writing process. For example, LLMs were used to structure the results chapters, catch overreach in conclusions against the data we report, identify awkward wording, and suggest where additional analysis might be needed. We further used LLMs in the writing process to develop our list of abbreviations at the end of the report, flag every page where an abbreviation was first mentioned but not spelled out, generate the initial figures in the report to provide a vision for our production team as the report moves toward publication, perform a full desk edit of our initial draft reports, and compile all our footnoted references into a reference list at the end of the report and format them.

At every stage, human researchers designed the research approach, developed the framework and variable definitions, set the decision rules, reviewed the outputs against the actual data, and made the final analytical judgments. The multi-agent architecture was designed so that no single model's output was accepted without review by either a second model from a different provider or a human analyst. The manual review process, conducted by a team of seven research analysts on a randomly selected sample, served as the primary validation check on the automated pipeline and informed iterative improvements throughout. We conducted the descriptive statistics and all quantitative analyses using standard analytical tools. The analytical findings, interpretations, and all text in this report reflect our judgment, informed—but not determined—by AI-generated outputs.

We adopted this approach because the scale of the research—1,179 products, each evaluated across dozens of features with supporting evidence and rationale—exceeded what traditional, manual methods could accomplish within the project's timeline and resources. The use of AI tools enabled a breadth of coverage that would otherwise have required a substantially larger research team or a substantially narrower scope. The trade-off is that the dataset carries the known limitations of LLM-based research, which are addressed in the next section.

Limitations

The research approach carried meaningful limitations that readers should weigh when interpreting the findings.

Coverage and Representativeness

We used complementary search strategies so that products missed by any one strategy would have multiple opportunities to surface through another. Still, similar to any landscape assessment, this approach cannot guarantee complete coverage of every product that exists; products with limited public visibility or that use terminology outside our search strategies may not be represented. Products that did not appear in disaster and EM conference technology company lists, LLM-based web searches, commercial company index searches, targeted web scraping, or stakeholder and expert referrals are not in the landscape. Findings about proportions, patterns, and concentrations describe the landscape of the products that we identified, but generalizations beyond that set should be treated with appropriate caution.

The analysis is based almost entirely on publicly available information. Technology company websites, product documentation, press releases, procurement records, cloud marketplace listings, conference

materials, news coverage, patent filings, and regulatory records were the primary inputs. Technology companies with limited public presence are underrepresented, and characterizations that would require nonpublic information (e.g., contract terms, internal deployment details, or usage data) were not feasible at scale. Unlike more mature product categories, for which standardized consumer information exists (e.g., nutrition labels, feature sheets, independent ratings), the AI-enabled product space lacks comparable infrastructure. Technology companies do not supply standardized information; critical details are often not publicly available, and no standard-setting, regulating, or consumer-focused entity currently evaluates or tracks this information systematically. This information environment shaped what could be characterized at scale and what could not, which is itself a finding of the research that we discuss in the report.

The research was completed between October 2025 and March 2026. The AI-enabled product sector is moving rapidly, and information current at the time of data collection may already be outdated. Products have been updated, new products have entered the market, and technology company and product positions on key dimensions may have shifted since the research closed.

Readers should treat the landscape findings as descriptive of a substantial and broadly representative slice of the market with strong public presence, while recognizing that products with limited public visibility may not appear in the dataset.

LLM-Based Collection and Its Implications

One of the five product-identification strategies used LLM APIs to search the web for AI-enabled products against each task area. LLM capabilities are shaped by their training data, which has cutoff dates and uneven coverage across topics. Products with strong online presence and clear AI-related language were more likely to surface through LLM-assisted search than products that exist but are less documented or use different terminology. This implies that the actual market for these products is likely a much vaster landscape, the scale of which would still necessitate the use of LLMs given the project's timeline. The four other strategies were designed to complement LLM-assisted search for exactly this reason, but the LLM component carries its own selection effect that shaped which products surfaced and which did not.

Product feature extraction and initial verification were also performed by LLMs conducting web searches and interpreting technology company–published content. LLMs can misinterpret ambiguous language, conflate similarly named products or companies, and generate plausible-sounding but inaccurate information. The multistage verification process—using a different model from a different provider for verification, followed by manual review—was designed to mitigate these risks, but it cannot eliminate them. Technology company marketing language may influence LLM interpretation, and products described in aspirational terms may be classified differently than products described in precise technical language, even when the underlying capabilities are similar. Some product features required inferential judgments when information was ambiguous or incomplete; these inferences were documented in rationale fields but carry inherently higher uncertainty than classifications based on explicit technology company statements. The extraction and verification models were selected from different providers to reduce correlated biases, but both share broadly similar training data distributions and may exhibit correlated blind spots on specialized product categories.

Products with strong online presence and clear AI-related language were more likely to surface through LLM-assisted search. This implies that the actual market for these products is likely a much vaster landscape.

Some product characterizations rest on inference from available evidence rather than direct confirmation from technology companies. Connectivity requirements, deployment

scale, and AI governance claims, among others, were sometimes inferred from product documentation or architecture rather than confirmed through technology company–provided sources. We applied systematic criteria for these inferences—documented in the framework definitions—but readers should understand that specific product characterizations may require verification with the technology company before a procurement decision.

Validation Scope and Reliability

A team of seven research analysts manually reviewed a randomly selected sample of 240 products. The sample size was calculated at the product level for 90 percent confidence with a ≤10 percent margin of error relative to the full dataset. Because each product was evaluated across 30 features, the review encompassed approximately 7,200 individual field-level judgments.

The review was designed as a plausibility assessment: Reviewers evaluated whether each automated classification was reasonable given the publicly available evidence, using the same source URLs and rationale fields produced by the automated process. The review confirms the pipeline's interpretive consistency with human judgment but does not independently verify facts as reported by technology companies.

Percentage agreement across reviewers reached 93 percent. We note that percentage agreement can overstate reliability when category distributions are imbalanced, as is the case for several features in this dataset. Across all field-level judgments in the sample, each product feature achieved at least a 90 percent confirmation rate. This validation confirmed that the automated pipeline was producing results broadly consistent with human expert judgment and identified systematic issues that were corrected through iterative refinement. It was not designed to be—and should not be read as—a precise audit of the accuracy of any single variable. To mitigate remaining risk, we used LLMs to identify patterned inconsistencies in the data, including a deliberate exploration of issues (e.g., the evidence did not support the categorization of a product on a certain dimension). When we identified a systemic issue with the underlying data, we corrected it, thereby increasing our confidence in the data.

Reviewers noted during the manual review process that some product features were difficult to classify even with human judgment, particularly when technology company–published information was vague, inconsistent across sources, or used nonstandard terminology—context that applies equally to the automated process. The technical expertise of the review team may also mean that the review standard reflects a level of AI literacy that a typical emergency manager or procurement officer may not share, and information that our reviewers could reasonably interpret from indirect evidence may be less accessible to potential buyers evaluating these products without a comparable technical background.

The qualitative analysis of rationale fields was LLM-assisted, meaning that the tool used to identify patterns shares the same class of limitations as the tool that generated the original data. Human researchers reviewed and validated the qualitative findings, but the initial pattern identification was not fully independent of the extraction process.

Aggregate findings about market-level patterns rest on a sound validation foundation. Specific characterizations of any individual product, particularly on dimensions for which evidence was inferred from indirect signals, may benefit from confirmation with the technology company before procurement decisions are made.

Definitional Choices and Analytical Methods

Definitional choices were made deliberately to err in defensible directions: Privacy concerns defaulted to "Yes" when inherently sensitive data types were present (i.e., erring toward caution); legal concerns defaulted to "No" unless substantial evidence existed (i.e., erring against unfounded attribution). The 45 EM task areas were developed from authoritative EM documents and validated through expert review. Alternative taxono-

mies or a different granularity, such as those that reflect only the primary work conducted by emergency managers as opposed to the breadth of their work, would produce different coverage distributions. The classification system of whether a product was purpose-built for EM reflects in part how explicitly a technology company markets to the EM community, which may undercount products with strong functional relevance but no EM-specific marketing. Readers comparing prevalence rates across categories should keep these defaults in mind: Privacy figures reflect a precautionary standard, and legal figures reflect a conservative standard.

Scope of the Assessment

The scope of 1,179 products and 45 task areas from 717 technology companies was a deliberate choice to characterize the full breadth of the landscape rather than evaluate a small set of products in depth. In this report, we characterize what products exist, how they are structured, and what conditions affect their adoption readiness. We did not directly observe what products are being used successfully in specific SLTT jurisdictions and nonprofits or how those products perform relative to the task area they could support. Adoption in practice requires fieldwork beyond the scope of this assessment, and the report should be read as a landscape characterization rather than a catalog of current adoption.

The scale of the landscape required that breadth come at the cost of depth on any given product. Readers seeking a deep comparative evaluation of specific products should treat this report as a starting point rather than a substitute for that work.

Cost Estimation

Cost was a dimension in which these limitations compounded most visibly and where the cumulative research investment was greatest relative to what could ultimately be reported. As we described above, we pursued cost data through multiple approaches over the course of the project—from early manual searches of technology company websites, to the exploration of private-sector data sources, to sustained manual research of public procurement records, to a purpose-built automated pipeline—and none produced data we could have confidence in. The barriers were consistent across every approach: Pricing opacity, structural variation in access models, and the absence of standardized cost disclosure meant that even the most resource-intensive methods could not produce consistent or comparable cost figures across the market. The resulting estimates were insufficiently reliable for quantitative use. The pricing landscape is described qualitatively; the depth of cost analysis possible from public sources was limited by the market's practices, not by research effort.

APPENDIX C

Response Task Areas and AI-Related Patterns

In Table C.1, we list each response task area and provide a description of the task area, the total number of products we identified that are associated with the task area, the dominant AI techniques that support these products, and the patterns regarding what AI enables in the task area products.

TABLE C.1

Response Task Areas and AI-Related Patterns

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Situational Awareness, Hazard Monitoring, and Information Sharing	Tasks in this area are to continuously collect, verify, and synthesize incident status, impacts, and resource posture and maintain and share a common picture.	516	• Nearly universal ML and predictive modeling • CV is the second-most common (satellite, drone, camera networks) • NLP is prominent in OSINT and event detection • Multimodal AI appears in sensor fusion and multisource analysis platforms	AI is collecting and analyzing data at scale. It tracks changes and looks for patterns across datasets.
Public Information, Alerts, and Protective Guidance	Tasks in this area are to craft, issue, and disseminate time-critical alerts and ongoing public updates with clear protective actions.	131	• ML and predictive modeling (threat detection that triggers alerts) • NLP (message generation, translation) • Generative AI (drafting and optimization)	The dominant AI contribution is speed and reach amplification—AI does not decide what to say but makes message creation faster (e.g., generative AI drafting), broader (translation), and more targeted (geo-fencing and automated triggers).
Protective Action Decision and Implementation	Tasks in this area are to decide and carry out life-safety measures (e.g., evacuate, shelter-in-place), including zones, timing, and special-population support.	37	• ML and predictive modeling (fire spread, flood, traffic modeling) • CV (wildfire detection triggering decisions) • Planning, scheduling, and optimization (evacuation sequencing)	The AI helps decisionmakers model consequences and the timing of protective actions, not just execute them.
Evacuation and Transportation Operations	Tasks in this area are to plan, route, and move people to safety and coordinate transportation assistance and priority movement.	59	• ML and predictive modeling (traffic prediction, clearance estimation) • Planning, scheduling, and optimization (routing, signal control) • CV (vehicle and pedestrian detection at intersections)	AI's role is predominantly predictive modeling and optimization—predicting how evacuations will unfold, optimizing signal timing, and providing real-time feedback on corridor performance.

Table C.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
OnScene Security, Access, and Perimeter Control	Tasks in this area are to manage access restrictions, site security, road closures, and reentry to protect responders and the public.	48	• CV (dominant—license plate recognition, facial recognition, weapon detection, video analytics) • ML and predictive modeling (alert scoring, behavior analysis)	AI's role is automated detection and alerting at access control points.
Search and Rescue	Tasks in this area are to locate, access, stabilize, and remove people from dangerous environments to safe locations.	55	• CV (person and object detection) • Robotics and embodied AI (autonomous navigation, obstacle avoidance) • ML (target tracking, thermal signature classification) • Planning, scheduling, and optimization (mission planning)	AI enables faster, safer search by putting sensors where humans cannot go (or cannot go quickly). The dominant pattern is CV-based detection on autonomous or semiautonomous aerial platforms.
Public Health, Health Care, and EMS	Tasks in this area are to field triage and treatment, patient movement, health care coordination and surge, and immediate public-health actions.	67	• ML and predictive modeling (clinical prediction, surge forecasting) • CV (medical imaging, X-ray AI) • NLP (voice-to-text documentation, translation) • Generative AI (ambient clinical documentation, such as Epic AI Charting)	AI in this task area follows two tracks: 1. clinical decision support (e.g., stroke imaging, sepsis detection, patient flow) that targets hospitals and healthcare systems 2. field operations support (e.g., EMS documentation, patient tracking, dispatch protocols, responder monitoring) that targets pre-hospital providers.
Fatality Management Services	Tasks in this area are to recover, identify, track, and care for the deceased with dignity and support family notifications.	10	• ML and predictive modeling (biometric matching, DNA analysis algorithms) • Knowledge representation (genealogical reasoning)	AI in fatality management is concentrated in biometric matching at scale—processing degraded DNA, degraded fingerprints, and facial images against large databases.
Fire Management and Suppression	Tasks in this area are to control and suppress fires to protect life, property, and the environment.	45	• CV (smoke and fire detection, satellite analysis) • ML and predictive modeling (fire spread simulation, risk scoring) • Planning, scheduling, and optimization (suppression mission planning) • Robotics (autonomous aircraft)	Fire management has the most end-to-end AI pipeline of any single-hazard task area the AI helps with, going from early detection to spread prediction to suppression planning to autonomous suppression to post-fire assessment.

Table C.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Environmental Hazard Response and Responder Health and Safety	Tasks in this area are to detect, contain, and mitigate acute environmental and hazardous material threats and monitor and protect responder and public health and safety.	44	• ML and predictive modeling (physiological prediction, gas identification algorithms) • Robotics (hazardous area entry) • CV (environmental sensing) • Knowledge representation (HazMasterG3 rules-based reasoning)	AI serves two distinct functions: 1. hazard detection—identifying what the threat is (chemical identification, gas detection, water contamination) 2. keeping humans out of harm's way—responder biometric monitoring, robotic zone entry, and safety compliance monitoring.
Shelter, Mass Care, and Human Services	Tasks in this area are to provide emergency shelter, food, hydration, sanitation, supplies distribution, and short-term support services.	25	• ML and predictive modeling (needs matching, demand forecasting) • NLP (chatbots, intake automation) • Knowledge representation (resource matching rules) • Planning, scheduling, and optimization (supply distribution)	This task area has thinner AI coverage than most operational task areas. The strongest products here are chatbots providing information access and donation-matching platforms.
Family Reunification and Welfare Checks	Tasks in this area are to receive and resolve missing-person reports, conduct welfare checks, and support safe reunification.	18	• ML (biometric matching, crash detection) • CV (facial recognition) • NLP (voice transcription/translation)	Products bifurcate into forensic identification (primarily for deceased and unidentified persons, overlapping with Fatality Management) and consumer safety apps (for families seeking to confirm their loved ones' safety).
Volunteer, Donations, and Offers Management	Tasks in this area are to solicit, vet, match, and coordinate volunteers and in-kind and financial donations to documented needs.	19	• Planning, scheduling, and optimization (volunteer scheduling) • ML (donor prediction, skills matching) • NLP (communication automation) • Generative AI (volunteer page creation, campaign content)	The dominant AI archetype is intelligent matching—connecting supply (volunteers, goods, and donations) to demand (documented needs).
Assistance Requests Intake, Triage, and Dispatch	Tasks in this area are to receive, verify, and prioritize requests for help and dispatch or refer appropriate resources.	96	• NLP (transcription, translation, conversational AI) • ML (call triage, priority scoring, routing) • Generative AI (call summarization, non-emergency agent)	The dominant innovation is AI as copilot for telecommunicators—not replacing humans but augmenting them with transcription, translation, and summarization so that they can process more calls with less cognitive load.

Table C.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Resource Request, Allocation, and Logistics	Tasks in this area are to request, prioritize, stage, track, and deliver personnel, equipment, and commodities to meet operational needs.	118	• ML and predictive modeling (ETAs, demand forecasting, equipment health) • Planning, scheduling, and optimization (routing, crew assignment, warehouse operations) • CV (dashboard camera safety, condition monitoring) • Robotics (autonomous delivery)	AI is strongest in visibility (knowing where things are) and optimization (deciding where to send them). Most products in this space are commercial fleet and supply chain products adapted for emergency use rather than purpose-built for disaster logistics. Drone delivery represents an emerging autonomous logistics capability.
Emergency Debris Clearance and Route Opening	Tasks in this area are to remove immediate life-safety hazards and open priority routes to enable rescue and aid delivery.	9	• CV (damage detection) • ML (route assessment) • Robotics (autonomous equipment operation)	This is a physically intensive task area in which AI's role is emerging but limited. The autonomous heavy equipment cluster represents the highest-autonomy AI application in debris operations—a single operator controlling multiple machines. AI's strongest contribution is in identifying where debris is (satellite and road imagery) rather than broadly tracking, managing, and removing it.
Critical Infrastructure and Essential Services Stabilization (Immediate)	Tasks in this area are to take immediate actions that stabilize essential services (e.g., temporary power and fuel, water and wastewater workaround) to protect life and property and support operations.	76	• ML and predictive modeling (dominant—outage prediction, asset health, demand forecasting) • Planning, scheduling, and optimization (distributed energy resource dispatch, grid balancing) • CV (sewer and pipeline inspection) • Knowledge representation (rule-based grid control)	AI's highest-value application is in predictive outage modeling—converting weather forecasts into expected outage counts, locations, and crew requirements days in advance. This directly supports pre-staging decisions. The microgrid and virtual power plant cluster represents AI-enabled islanded resilience—keeping critical facilities powered when the main grid fails. Water and sewer products focus on rapid damage assessment and workaround identification.

Table C.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Communications and IT Continuity for Response Operations	Tasks in this area are to maintain and restore mission-critical communications and basic IT needed to coordinate and sustain response.	68	• ML and predictive modeling (anomaly detection, threat detection, network optimization) • NLP (transcription, translation, noise cancellation) • Generative AI (incident summarization, root-cause analysis)	AI in this task area serves a meta-function—it keeps the technological infrastructure that all other task areas depend on operational. AI operations products reduce alert fatigue (correlating thousands of alerts into actionable incidents), cybersecurity AI prevents disruption, and unified communications-as-a-service AI (noise cancellation, transcription) improves communication quality in noisy field environments.
Incident Coordination and Operational Decision Support	Tasks in this area are to coordinate multi-organization actions, set objectives, manage operational periods, and document decisions and priorities.	340	• ML and predictive modeling (nearly universal) • NLP (summarization, conversational queries) • Generative AI (briefing generation, decision support) • Knowledge representation (ICS workflows, protocol logic) • Planning, scheduling, and optimization	AI plays two distinct roles: 1. acceleration—getting information to decisionmakers faster through summarization, fusion, and autonomous scouting 2. prediction—modeling scenarios to inform resource positioning and operational planning. Products span from platforms that structure the entire incident management life cycle to point solutions providing specific decision inputs (e.g., weather, fire models). The emerging trend is generative AI assistants embedded in EOC platforms that can summarize board data, generate situation reports, and answer ad hoc queries in natural language. This represents AI moving from providing data to providing synthesized judgment support.

Table C.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Initial Damage and Impact Assessment	Tasks in this area are to rapidly assess damages and priority needs, confirm the cessation of immediate threats, and inform next steps.	209	• CV (dominant—the core technology for visual damage assessment) • ML and predictive modeling (damage classification, risk scoring) • Multimodal AI (fusing optical, SAR, and sensor data) • Generative AI (report generation from field data)	The dominant pattern is automated image analysis at scale—turning satellite, aerial, or ground-level imagery into damage classifications faster than manual inspection. The pipeline typically flows from imagery capture to AI classification to GIS visualization to FEMA-aligned reporting.

NOTE: ETA = estimated time of arrival.

APPENDIX D

Disaster Recovery Task Areas and AI-Related Patterns

In Table D.1, we list each recovery task area and provide a description of the task area, the total number of products we identified that are associated with the task area, the dominant AI techniques that support these products, and the patterns regarding what AI enables in the task area products.

TABLE D.1

Disaster Recovery Task Areas and AI-Related Patterns

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Damage and Impact Assessment	Tasks in this area are to identify and quantify physical, social, and economic impacts after an incident. They cover rapid building damage surveys, infrastructure status checks, community needs assessments, and mapping of affected populations and critical facilities. This task area supports the prioritization of life-safety risks and time-sensitive recovery actions.	260	• CV (property and infrastructure inspection) • ML (damage scoring, economic modeling) • Multimodal (fusing imagery, sensor, and demographic data)	The dominant pattern is automated image analysis at scale—turning satellite, aerial, or ground-level imagery into damage classifications faster than manual inspection. The recovery version emphasizes comprehensiveness and accuracy (i.e., every property, structural engineering detail) over speed (i.e., rapid initial picture). Products shift from satellite-scale to property-level inspection and from binary damaged-or-not assessments to detailed condition reports with repair estimates.
Infrastructure and Lifelines Restoration	Tasks in this area are to restore critical lifelines, such as power, water, wastewater, transportation, and communications, so that communities can function safely. They cover prioritizing repairs, sequencing restoration, coordinating field crews, issuing boil-water or outage advisories, and tracking system status.	71	• ML and predictive modeling (dominant—outage prediction, asset health, demand forecasting) • Planning, scheduling, and optimization (crew dispatch, restoration sequencing) • CV (condition assessment) • Knowledge representation (asset life-cycle rules, fault diagnosis)	The strongest AI application is in electric utility restoration—a mature market with dedicated products for outage prediction, crew optimization, and customer communication. Water and wastewater and telecommunications restoration have fewer purpose-built AI products. The core AI value proposition is reducing restoration time by optimizing the task sequence: identify damage, prioritize repairs, dispatch crews, and track progress.
Housing and Shelter Recovery	Tasks in this area are to transition survivors from emergency shelter to safe and accessible interim and long-term housing. They cover assessing housing needs, matching survivors to available units, siting temporary housing, tracking habitability and reentry status, and coordinating renovations or rebuilds.	30	• ML (housing matching, risk scoring) • CV (property inspection, building plan review) • NLP (case management, chatbots) • Generative AI (code compliance checking)	AI's strongest application is in accelerating the rebuild pipeline—from code compliance checking (e.g., reducing permit review from weeks to days) to construction management (e.g., optimizing rebuild schedules).

Table D.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Health, Education, Behavioral Health, and Social Services Recovery	Tasks in this area are to restore and strengthen the systems that keep people physically healthy, mentally well, socially supported, and connected to education, especially for those who are directly or indirectly affected by disaster. Tasks cover assessing damage and the functionality of hospitals, clinics, pharmacies, behavioral health facilities, schools, child care centers, and social services offices; reestablishing essential medical care (e.g., primary care, chronic disease management, maternal and child health), including temporary or mobile services for facilities that are damaged; expanding and coordinating behavioral health and psychosocial support, including crisis counseling, trauma-informed care, and support groups; restoring kindergarten through grade 12 schools, higher education, and child care services, including temporary learning environments, transportation adjustments, and accommodations for displaced students and staff; ensuring the continuity of critical social services (e.g., food assistance, disability services, child welfare, services for older adults and people with access and functional needs); and identifying and prioritizing high-risk and disproportionately affected populations and ensuring that they can access health, education, and social supports during recovery.	39	• NLP (crisis text analysis, behavioral monitoring, documentation) • ML (risk scoring, triage, activity pattern detection) • Generative AI (conversational wellness companions) • CV (student content monitoring, video assessment)	The strongest AI contributions are in 1. behavioral health screening at scale—identifying who needs help before they ask for it (student monitoring, crisis text triage, video-based assessments) 2. social determinants of health analytics—mapping where social vulnerability intersects with disaster impact to target recovery resources.
Individual and Household Assistance and Case Management	Tasks in this area are to help individuals and households navigate assistance options and recover basic stability. They cover intake and triage, eligibility screening, referral to programs (housing, food, financial help), follow-up, and coordination among multiple providers around a single survivor or family.	64	• NLP (conversational intake, document processing, multi-lingual support) • ML (eligibility prediction, risk scoring, case prioritization) • Generative AI (response drafting, form generation) • Knowledge representation (program rules, eligibility logic)	The dominant AI archetype is navigation and matching. Benefits navigation is the fastest-growing segment.

Table D.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Economic and Workforce Recovery	Tasks in this area are to stabilize and revitalize local economies and employment. They cover business impact assessments, small business support, coordination of reopening plans, workforce retraining and placement, and the prioritization of investments that protect jobs and community anchors.	41	• ML and predictive modeling (skills matching, labor forecasting, lending risk) • NLP (resume and skill extraction, learning content) • Generative AI (career coaching, learning assistance) • Knowledge representation (skill taxonomies, program eligibility)	AI's primary role is accelerating the matching of people to economic opportunities—whether they are jobs, training, or capital. Skills-based matching (i.e., ignoring degree requirements) is a dominant theme, directly relevant when disaster-displaced workers need rapid reemployment in unfamiliar sectors.
Debris, Waste, and Environmental Safety Management	Tasks in this area are to remove debris and hazardous materials to restore safe access and reduce secondary risks. They cover debris volume estimation, route planning and staging, prioritization of clearance for critical facilities and neighborhoods, environmental testing, and coordinating disposal and recycling.	22	• CV (debris detection, waste sorting, environmental sensing) • ML (route optimization, compliance) • Knowledge representation (waste classification rules, Resource Conservation and Recovery Act compliance)	The most effective AI contribution is in automated detection (where debris is) and logistics optimization (how to route crews and trucks). The products that exist focus on waste tracking and compliance rather than the massive logistics challenge of debris estimation, staging, hauling, and monitoring. AMP's robotic sorting is innovative but designed for recycling facilities, not disaster debris sites.
Natural and Cultural Resources Restoration	Tasks in this area are to assess and restore damaged natural systems (e.g., coasts, forests, rivers) and cultural or historic sites. They cover impact assessments, the prioritization of restoration projects, permitting coordination, and engagement with communities regarding culturally significant sites and practices.	10	• CV (dominant—satellite image analysis, species identification) • ML (change detection, carbon estimation, predictive ecology) • Robotics (forestry automation)	AI enables continuous, large-scale environmental monitoring that would be impossible manually, such as tracking millions of acres of forest, identifying species from acoustic recordings, or mapping coastal erosion from space. The products are primarily monitoring and assessment products rather than restoration execution products.

Table D.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Community Engagement, Communication, and Information Access	Tasks in this area are to ensure survivors and stakeholders receive timely, accessible, and trustworthy information and can provide input. They cover recovery public messaging, community meetings, two-way feedback, rumor control, translation and interpretation, and accessible formats for people with disabilities and limited English proficiency.	80	• NLP (dominant—translation, transcription, chatbot conversation, sentiment analysis) • Generative AI (content generation, message drafting) • ML (sentiment scoring, need prediction)	The core challenge is communicating complex recovery information to diverse populations, and AI addresses this through 1. translation and interpretation across dozens of languages 2. chatbots that provide instant answers without requiring call center staff 3. sentiment analysis to track whether communications are working. The two-way feedback component is addressed by engagement platforms that analyze open-ended responses.
Resource, Logistics, and Supply Chain Management for Recovery	Tasks in this area are to obtain, allocate, and track resources needed for recovery operations, including materials, equipment, and personnel. They cover identifying gaps, prioritizing requests, staging supplies, coordinating with technology companies, and tracking deliveries and usage across sites.	41	• ML (predictive ETAs, demand forecasting, risk scoring) • Planning, scheduling, and optimization (routing, warehouse operations, construction scheduling) • Knowledge representation (supply chain graph models) • NLP (disruption event detection from news)	The strongest AI capability is in disruption detection and adaptive replanning. When a supplier fails or a route closes, AI identifies alternatives.
Volunteer, Donations, and Community Support Coordination	Tasks in this area are to organize spontaneous and affiliated volunteers, manage in-kind and financial donations, and align community support with real needs. They cover volunteer registration and assignment, donations intake and matching, avoiding duplication of effort, and capturing community capacity for recovery.	18	• ML (skills matching, donor propensity) • NLP (communication automation, campaign content) • Generative AI (volunteer opportunity creation, fundraising content) • Planning, scheduling, and optimization (assignment, scheduling)	The recovery volunteer and donation challenge differs from response: It requires sustained coordination over months rather than surge mobilization. AI helps by matching skills to specialized recovery task areas (e.g., mold remediation, case management, construction) and maintaining donor engagement over longer timelines.

Table D.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Recovery Governance, Planning, and Interagency Coordination	Tasks in this area are to structure how recovery decisions are made and coordinated across agencies and organizations. They cover setting up recovery task area forces, developing and updating recovery plans, tracking progress against objectives, resolving conflicts, and aligning local, state, and federal roles.	35	• Generative AI (briefing generation, plan drafting, meeting summarization) • NLP (document analysis, natural-language querying) • ML (progress tracking, scenario modeling)	AI's role in governance is primarily about information synthesis and decision acceleration—turning large volumes of recovery data into concise briefings, tracking corrective actions, and enabling leaders to ask questions in natural language rather than running reports. The generative AI wave is particularly effective here because it addresses the enormous documentation burden of recovery governance (e.g., plans, reports, meeting minutes, compliance documentation).
Financial Recovery, Grants, Cost Tracking, and Insurance	Tasks in this area are to enable institutions, individuals, households, and businesses to secure, manage, and account for the money needed to recover and do so in a way that meets eligibility, compliance, and audit requirements. For public agencies and NGOs, they cover identifying and prioritizing funding sources (e.g., federal, state, local, philanthropic, insurance) for recovery projects; preparing, submitting, and managing grant, loan, and reimbursement applications (e.g., public assistance, mitigation funding); cost tracking and documentation (e.g., labor, equipment, contracts, materials) to support reimbursement and avoid de-obligation; managing compliance and reporting requirements across multiple funders and programs; and aligning and braiding multiple funding streams to support priority recovery projects. For individuals, households, and businesses, they cover understanding and applying for financial assistance programs (e.g., grants, loans, cash assistance, tax relief); contacting and working with insurance providers (e.g., home, renters, flood, business, crop), filing claims, tracking status, documenting losses (e.g., photos, receipts, estimates) and navigating appeals and deadlines; and getting connected to financial counseling, legal aid, and small-business support focused on recovery.	97	• ML (risk scoring, fraud detection, claim prediction) • CV (damage estimation from imagery, document OCR) • NLP (grant proposal drafting, document extraction) • Generative AI (grant writing, claims summarization) • Knowledge representation (compliance rules, eligibility logic)	The fundamental AI value proposition is accelerating money flow—faster claims processing, faster grant applications, and faster cost documentation.

Table D.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Land Use, Rebuilding, and Mitigation Planning Through Recovery	Tasks in this area are to guide where and how rebuilding occurs to reduce future risk (build back better). They cover post-disaster land use decisions, buyout or relocation programs, elevation and retrofitting plans, updates to building codes, and the integration of mitigation into recovery projects.	24	• CV (plan review, building detection) • NLP (code interpretation) • Generative AI (code Q&A, plan analysis) • Knowledge representation (code compliance rules) • ML (climate projection, risk scoring)	AI's strongest contribution is in accelerating the permitting pipeline—the bottleneck that slows rebuilding after disasters. Automated plan review can reduce review times from weeks to hours. Climate risk models inform where rebuilding should occur (and where it should not), supporting buyout and relocation decisions. The connection between rebuilding and mitigation (e.g., stronger codes, elevation requirements) is addressed by code compliance products that enforce updated standards.
Data, Information Management, and Situational Awareness for Recovery	Tasks in this area are to collect, integrate, analyze, and share recovery-relevant data. They cover maintaining common operating pictures, integrating geospatial and sectoral data, tracking key recovery indicators, and providing decision support dashboards or reports to leaders and the public.	179	• ML (analytics, pattern detection) • Generative AI (summarization, conversational data access) • NLP (document processing, search) • CV (imagery analysis, video analytics)	AI's role is making recovery data accessible and actionable, converting raw data into dashboards, enabling leaders to ask questions in natural language, and automatically tracking recovery metrics. The generative AI wave has significantly affected this space through conversational analytics that let nontechnical users interact with data.

Table D.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Private-Sector and Public-Private Partnership Recovery Coordination	Government, private sector, and nongovernmental partners work together in structured or semi-structured ways (e.g., P3s, BEOCs, sector coalitions) to coordinate recovery actions, particularly around lifelines, critical services, and economic stability. They cover establishing, staffing, and operating public-private structures for recovery (e.g., P3s, BEOCs, sector working groups); creating shared situational awareness about impacts to critical businesses, infrastructure, supply chains, and the workforce; coordinating the restoration of community lifelines and essential services (e.g., power, telecommunications, fuel, grocery and food retail, pharmacies, banking, logistics); jointly identifying and addressing supply chain disruptions (e.g., fuel, medical supplies, construction materials) that are affecting recovery; aligning business reopening, relocation, or reinvestment decisions with community recovery priorities and land-use or mitigation goals; and facilitating information sharing, mutual aid, and resource sharing between private entities, NGOs, and the government (e.g., sharing generators, facilities, data, or expertise).	13	• ML (risk scoring, disruption detection) • Knowledge representation (supply chain mapping) • NLP (event monitoring, intelligence sharing)	AI products provide information inputs to private sector and P3s rather than managing the partnerships.

NOTE: BEOC = business emergency operations centers; NGO = nongovernmental organization; P3 = public-private partnership.

APPENDIX E

Preparedness Task Areas and AI-Related Patterns

In Table E.1, we list each preparedness task area and provide a description of the task area, the total number of products we identified that are associated with the task area, the dominant AI techniques that support these products, and the patterns regarding what AI enables in the task area products.

TABLE E.1

Preparedness Task Areas and Patterns Regarding What Current AI Products Enable

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Program Governance, Authorities, and Administration	Tasks in this area are to establish the organizational backbone of preparedness. They cover the designation of an EM entity and leader; the authority and responsibility to coordinate mitigation, response, and recovery activities for the jurisdiction; the development of fiscal and administrative procedures that work before, during, and after disasters, including those procedures needed to request, receive, and manage funds to support mitigation, response, and recovery activities; and tracking and addressing legislative and regulatory changes that affect EM programs, task areas, and activities.	44	• Generative AI (document drafting, legislative analysis, policy Q&A) • NLP (regulatory monitoring, contract extraction, natural-language search) • ML (compliance scoring, risk classification) • Knowledge representation (rules engines, eligibility logic)	AI primarily reduces the administrative burden of maintaining an EM program: tracking regulatory changes, managing contracts, processing grants, and maintaining records. Generative AI assistants are rapidly entering this space, enabling staff to query policy documents and draft correspondence. The legislative tracking cluster is notable for predicting bill outcomes and summarizing regulatory impacts.
Hazard Identification and Vulnerability Assessment	Tasks in this area are to systematically identify the full scope of natural and human-caused threats and hazards that may affect the jurisdiction and assess who and what is vulnerable to those threats. They cover gathering and analyzing technical, spatial, demographic, land-use, economic, transportation, and environmental data to build a shared understanding of conditions within and across the jurisdiction; identifying hazards through multiple sources, such as historical records, scientific studies, and subject-matter expertise; profiling each hazard to understand its characteristics, frequency, magnitude, and geographic extent; assessing the vulnerability of people, property, critical infrastructure, the environment, and EM operations to identified hazards, with particular attention to populations with access and functional needs; identifying community strengths, issues, opportunities, and constraints that shape readiness; and determining the stakeholders—across the government, private sector, and whole community partners—who need to be engaged throughout the assessment process. The results provide the factual foundation on which risk assessment, capability analysis, and all subsequent preparedness planning activities are built.	44	• ML and predictive modeling (dominant—risk scoring, catastrophe modeling, climate projection) • CV (building detection, imagery analysis, property assessment) • Knowledge representation (hazard classification frameworks)	AI is fundamentally transforming hazard and vulnerability assessment by enabling property-level risk scoring at a national scale—what previously required field visits or coarse zone-based estimates. The insurance industry has driven much of this innovation, and government EM programs increasingly access the same data. The core AI technique is combining aerial and satellite imagery with ML models trained on historical loss data.

Table E.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Risk Assessment	Tasks in this area are to create and maintain a repeatable, jurisdiction-wide picture of risk by analyzing the likelihood and consequences of the threats and hazards identified through hazard identification and vulnerability assessment. They cover estimating the probability and potential impact of each hazard across multiple dimensions—public health and safety, responder safety, continuity of government services, critical infrastructure and key resources, community lifelines, economic conditions, and public confidence; establishing context-specific capability targets that define the level of capability the jurisdiction needs to manage those risks; mapping and analyzing critical supply chains, particularly those supporting lifeline-relevant goods and services, to identify vulnerabilities, interdependencies, and single points of failure; conducting outreach with supply chain partners to execute resilience actions and assess improvements; and linking supply chain risk findings to distribution management planning as a bridge to sustain communities until normal supply systems are restored. The risk assessment process drives prioritization across all preparedness activities—informing planning, training, exercise design, resource investment, and mutual aid strategies—and should be updated periodically as hazard conditions, community characteristics, and capability levels change.	54	• ML and predictive modeling (nearly universal—probability estimation, consequence modeling, risk scoring) • NLP (event monitoring, intelligence extraction) • Generative AI (risk summarization)	Risk assessment is where ML and predictive modeling are most purely expressed—the task is fundamentally about estimating probability and consequence. AI enables continuous, automated risk scoring that updates with new data and replaces periodic manual assessments. The supply chain risk cluster is notable for monitoring thousands of global data sources 24/7 to detect risks before they materialize. The connection between risk assessment and capability/resource planning (called for in the task area definition) is weakly automated: Risk scores are produced, but their translation into preparedness requirements is largely manual.
Capacity and Capability Analysis	Tasks in this area are to evaluate jurisdictional mitigation, response, and recovery capacity and capability against the results of the hazard identification, vulnerability assessment, and risk assessment processes. They also establish capacity and capability gaps. The results are used to drive preparedness efforts in the areas of planning and organizing, training, exercises, and resource management.	14	• ML (skills matching, gap analysis, resource optimization) • NLP (skills extraction, plan analysis) • Generative AI (capability report generation)	This is the most underserved task area across all of the 45 task areas in our framework. The gap between what the task area definition calls for—evaluating mitigation, response, and recovery capacity against assessed risks and establishing capability targets—and what products deliver is significant. Most products address workforce capability (skills inventories) rather than operational capability (Can we evacuate 100,000 people in 12 hours? Can we shelter 50,000? Do we have enough swift-water rescue teams?).

Table E.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Integrated Preparedness Planning and Plan Synchronization	Tasks in this area are to develop, integrate, and maintain the full suite of plans that a jurisdiction needs to mitigate, respond to, and recover from the threats and hazards identified through risk assessment. These plans should include an emergency operations plan, mitigation plan, and recovery plan, among others. Planning involves getting the buy-in of jurisdiction leadership into the planning process, mapping the community to identify stakeholders important to the planning effort, gathering data to inform a shared fact base to inform planning, the development of plans (including structures, processes, tactics, strategies and/or projects), synthesizing information gathered, identification of resources to support the plan, plan implementation and monitoring, and plan maintenance and update. Organizing encompasses the activities through which a jurisdiction establishes, staffs, and maintains the organizational structures, management systems, facilities, and administrative foundations needed to coordinate mitigation, response, and recovery.	30	• Generative AI (plan drafting, conversion) • NLP (document analysis, search) • Knowledge representation (plan structure and logic) • ML (scenario-based resource identification)	Generative AI is having the most transformative impact on this task area—converting legacy plans from static documents into structured data, generating plan components from prompts, and enabling natural-language interaction with plan content.
Resource Management, Mutual Aid, Logistics, Distribution, and Supply Chain Resilience	Tasks in this area are to prepare to manage personnel, teams, facilities, equipment, and supplies before, during, and after incidents by identifying and typing resources, credentialing personnel, planning for resource acquisition and inventory, and building mutual aid mechanisms to obtain resources the jurisdiction does not own. They cover defining mutual aid agreement provisions (e.g., reimbursement, licensure recognition, mobilization procedures, interoperability protocols, and resource packaging standards).	44	• ML (risk scoring, demand forecasting, credential verification) • Knowledge representation (supply chain graphs, compliance rules) • NLP (contract analysis, regulatory monitoring) • Planning, scheduling, and optimization (warehouse operations, workforce scheduling)	AI is strongest in the supply chain resilience component—monitoring suppliers, mapping dependencies, and detecting risks. Credentialing is emerging with AI verification.

Table E.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Community Engagement and Whole Community Partnerships	Tasks in this area are to build, maintain, and leverage relationships with the full scope of whole-community stakeholders—individuals, families, faith-based organizations, community-based organizations, nonprofit groups, advocacy organizations, schools and academia, businesses, and media outlets—to strengthen collective preparedness, ensure inclusive and equitable preparedness efforts, and expand the capacity that is available for mitigation, response, and recovery. They cover engaging faith-based and community organizations as partners by understanding their capabilities, integrating them into planning processes, and establishing formal coordination mechanisms; building and sustaining private-public partnerships beyond information sharing, including leveraging private sector resources, expertise, and infrastructure for mutual benefit; conducting public preparedness education campaigns that empower individuals and families to take protective actions; ensuring that preparedness activities address the access and functional needs of children, older adults, individuals with disabilities, people with limited English proficiency, and other underserved or underrepresented populations; integrating community stakeholders into planning processes to ensure that plans reflect actual community needs, assets, and priorities; and using the whole-community approach to broaden participation and ownership of preparedness beyond government agencies.	36	• NLP (dominant—translation, sentiment analysis, chatbot conversation) • Generative AI (content creation, video generation) • ML (feedback analysis, engagement optimization)	AI addresses the language and accessibility barriers that limit whole-community engagement. Real-time translation and chatbots (available 24/7 in multiple languages) extend reach far beyond traditional town halls. The “partnership building” component (engaging faith-based organizations, building public-private partnerships) remains a fundamentally human activity that AI supports informationally but does not drive.

Table E.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Training	Tasks in this area are to build and sustain the knowledge, skills, and abilities of personnel who have roles and responsibilities in mitigation, response, and recovery. They cover conducting training needs assessments, informed by the results of hazard identification, risk assessment, capability analysis, and after-action findings from exercises and real-world incidents; developing and managing a multiyear training plan that aligns training priorities with identified capability gaps and plan requirements across mitigation, response, and recovery functions; employing a progressive approach in which foundational courses build toward more-advanced and more-integrated training activities; delivering training on plans, procedures, systems, and equipment to personnel across all levels of government, the private sector, and whole-community partners, including training on mitigation strategies and products (such as hazard mitigation planning and risk reduction measures); training on response operations (including the operational coordination, communication, and resource management concepts established in NIMS); training on recovery processes (including damage assessment, disaster housing, continuity of services, and long-term community recovery planning); incorporating training on the access and functional needs of children, older adults, individuals with disabilities, people with limited English proficiency, and other underserved populations; tracking individual and organizational training completions and credential requirements to maintain qualification and readiness standards; using training outcomes to inform exercise design, update plans and procedures, and validate that personnel can execute their assigned responsibilities under realistic conditions.	37	• Generative AI (scenario generation, content creation, conversational coaching) • NLP (transcription, performance analysis, natural-language coaching) • ML (personalized learning paths, performance prediction) • Knowledge representation (protocol logic for dispatch training)	AI can generate unlimited, unique training scenarios on demand, replacing the constraint of finite prewritten scripts. AI dispatch simulators provide instant feedback, eliminating the need for dedicated human evaluators for every training session. Virtual reality and extended reality platforms create safe, repeatable environments for high-risk training. AI coaching assistants provide personalized guidance at scale.

Table E.1—Continued

Task Area	Description of the Task Area	Total Associated Products	Dominant Supporting AI Techniques	Patterns Regarding What Current AI Products Enable
Exercising and Continuous Improvement	Tasks in this area are to test and validate capabilities consistent with identified hazards, identify strengths and areas for improvement, and drive continuous improvement. They cover managing a multiyear exercise program; using a progressive approach, in which each exercise builds on previous exercises; ensuring broad public participation; developing exercise documentation; conducting exercise evaluation through trained evaluators who assess performance against objectives and core capabilities; performing after-action analysis, including trend and root cause analysis; documenting findings and corrective actions in AARs and improvement plans; tracking the implementation of corrective actions through completion; and using results to inform updates to plans, organizational structures, training priorities, and resource investments.	25	• Generative AI (scenario generation, AAR drafting) • NLP (performance analysis, transcription) • ML (quality scoring, pattern detection) • Knowledge representation (scenario logic, evaluation criteria)	Generative AI automates what was previously a labor-intensive, weeks-long manual process (e.g., scenario development, Master Scenario Events List creation, documentation). After-action analysis products automatically transcribe and analyze exercise communications.

NOTE: NIMS = National Incident Management System.

APPENDIX F

Examples of Observed Response, Recovery, and Preparedness Product Gaps

In Table F.1, we provide examples of gaps we observed in the products in the landscape by functional area and task area.

TABLE F.1
Examples of Observed Response, Recovery, and Preparedness Product Gaps

Functional Area	Task Area Label	Examples of Observed Gaps
Response	Situational Awareness, Hazard Monitoring, and Information Sharing	Despite enormous product density, relatively few products specialize in synthesizing information across domains into an integrated common operating picture. Most provide single-domain feeds (weather **or** satellite **or** social media). True multi-domain fusion remains concentrated in a handful of platforms. The task area definition emphasizes "maintain and share a common picture," which implies cross-domain synthesis that most products do not independently achieve.
Response	Public Information, Alerts, and Protective Guidance	Few products address protective guidance personalization—tailoring the recommended protective action to a recipient's specific situation (e.g., mobility limitations, location within a flood zone, access and functional needs). Most deliver uniform messages to geographic areas.
Response	Protective Action Decision and Implementation	Integration between hazard modeling (fire and flood) and evacuation management (zones and traffic) is handled by only a few technology companies. Most products cover only one side of the decision.
Response	Evacuation and Transportation Operations	Products primarily optimize vehicle movement rather than the broader evacuation operation (which includes notification, shelter activation, and special population transport). AI is strongest in predictive traffic modeling and real-time signal optimization. The transportation assistance component for access and functional-needs populations is underserved by AI products—no products specifically address the challenge of coordinating paratransit, ambulance, or accessible vehicle fleets during evacuation.
Response	On-Scene Security, Access, and Perimeter Control	The access and perimeter control function is overwhelmingly focused on detection and identification rather than physical access enforcement (e.g., doors, barriers), which remains largely hardware-driven.
Response	Search and Rescue	Few products address the stabilize component of search and rescue—providing medical assessment or treatment guidance at the point of rescue. The link between detection-by-drone and medical triage is largely manual.
Response	Fatality Management Services	AI products cluster almost entirely on identification. The logistical and compassionate dimensions of fatality management (e.g., morgue operations, family assistance centers, death certificate processing at scale) have minimal AI support.
Response	Shelter, Mass Care, and Human Services	Physical shelter operations (e.g., food, sanitation, supplies distribution) have limited AI penetration.

Table F.1—Continued

Functional Area	Task Area Label	Examples of Observed Gaps
Response	Family Reunification and Welfare Checks	The traditional welfare check function—dispatching someone to physically verify a person's status—has almost no dedicated AI support. This is a significant gap given the volume of welfare check requests after disasters.
Response and recovery	Volunteer, Donations, and Offers Management	When we addressed matching of volunteers to task areas, we found only one purpose-built disaster community coordination product to help the organizations that coordinate volunteers and donations work together. We had a similar finding related to donations management.
Response	Emergency Debris Clearance and Route Opening	Debris volume estimation, the prioritization of clearance routes to critical facilities, and the coordination of debris staging—all described in the task area definition—have minimal dedicated AI support. This is a significant gap given that debris removal is typically the largest single cost category in major disaster recovery operations.
Recovery	Housing and Shelter Recovery	Siting and managing temporary housing (FEMA mobile housing units, transitional units)—a major operational challenge—has virtually no dedicated AI support.
Recovery	Health, Education, Behavioral Health, and Social Services Recovery	After a major disaster, tens of thousands of people may need behavioral health support. The AI products available are primarily general-population wellness products or small-scale crisis services. There is virtually no AI product that specifically addresses the need to deploy and scale disaster-specific behavioral health services (e.g., psychological first aid, grief counseling, long-term trauma treatment) to a suddenly affected population at scale. AI-assisted triage (e.g., Crisis Text Line's AVA algorithm) is the closest capability, but it was not designed for post-disaster population-level deployment.
Recovery	Individual and Household Assistance and Case Management	The "coordination among multiple providers around a single survivor" function—true multiagency case coordination—remains weakly automated. Most case management platforms track a single organization's interactions rather than orchestrating across the full network of providers serving one household.
Recovery	Economic and Workforce Recovery	Small business–lending AI could significantly accelerate SBA disaster loan processing, but most products target general commercial lending rather than disaster-specific programs.
Recovery	Debris, Waste, and Environmental Safety Management	Debris volume estimation from satellite or aerial imagery, priority route clearance planning, contractor management, and debris monitoring site operations—core debris management functions—have almost no purpose-built AI products.
Recovery	Natural and Cultural Resources Restoration	The products are primarily monitoring and assessment products rather than restoration execution products. The connection to cultural resources (e.g., historic sites, archaeological preservation) is essentially absent from the product landscape.
Recovery	Resource, Logistics, and Supply Chain Management for Recovery	AI products overwhelmingly come from commercial supply chains and logistics and are adapted for recovery rather than being purpose-built. The strongest AI capability is in disruption detection and adaptive replanning—when a supplier fails or a route closes, AI identifies alternatives. Purpose-built humanitarian logistics products that go beyond disruption detection and adaptive replanning to also model disaster-specific constraints are rare (e.g., road closures, points of distribution operations, needs matching from social media).
Recovery	Volunteer, Donations, and Community Support Coordination	When we addressed matching of volunteers to task areas and tracking, we found only one purpose-built disaster community coordination product to help the organizations that coordinate volunteers and donations work together.

Table F.1—Continued

Functional Area	Task Area Label	Examples of Observed Gaps
Recovery	Land Use, Rebuilding, and Mitigation Planning Through Recovery	AI support for politically sensitive decisions—buyout programs, land-use changes, zoning updates—is limited to data provision rather than decision support.
Recovery	Private-Sector and Public-Private Partnership Recovery Coordination	This is one of the least developed task areas in terms of AI products. Most products are commercial supply chain or business continuity products rather than purpose-built partnership coordination platforms.
Preparedness	Risk Assessment	The connection between risk assessment and capability and resource planning (called for in the task area definition) is weakly automated—risk scores are produced, but their translation into preparedness requirements is largely manual.
Preparedness	Capacity and Capability Analysis	This is the most underserved task area across all 45 task areas. The gap between what the task area definition calls for (i.e., evaluating mitigation, response, and recovery capacity against assessed risks and establishing capability targets) and what the products deliver is significant. Most products address workforce capability (e.g., skills inventories) rather than operational capability (e.g., Can 100,000 people be evacuated in 12 hours? Can 50,000 people be sheltered? Are there enough swift-water rescue teams?). The explicit linking of capacity gaps to preparedness efforts (e.g., planning, training, exercises) is essentially unautomated.
Preparedness	Integrated Preparedness Planning and Plan Synchronization	The synchronization aspect (e.g., ensuring that emergency operations plans, continuity of operations plans, hazard-specific annexes, mutual aid agreements, and recovery plans are consistent) remains primarily a human coordination task area, with AI providing search and comparison support.
Preparedness	Resource Management, Mutual Aid, Logistics, Distribution, and Supply Chain Resilience	The traditional EM resource management functions (e.g., NIMS resource typing, mutual aid activation, resource ordering and tracking) have less AI penetration than commercial supply chain management. Credentialing is emerging with AI verification, but the licensure recognition and interoperability protocol challenges described in the task area remain primarily procedural.
Preparedness	Exercising and Continuous Improvement	The task area definition calls for a multiyear exercise program with strategic sequencing. No AI products address exercise program strategy—what to exercise, in what order, and with what objectives based on assessed risks and capability gaps.

APPENDIX G

Sequenced Action Agendas for Stakeholder Groups to Support Adoption and Diffusion

In this appendix, we provide sequenced action agendas for the following stakeholders: emergency managers and others who support EM functions, technology companies, policymakers, and researchers.

An Agenda for Emergency Managers and Others Who Support EM Functions

Much of what determines whether and when emergency managers and others who support the work of EM will be funded and allowed to access and use AI-enabled products sits outside this group's control. However, what this group does control is still meaningful. Personnel can prepare so that when conditions allow, they are ready to implement, sustain, and standardize AI-enabled products. Table G.1 lays out a sequenced agenda across three time horizons for personnel in EM offices and the organizations that support EM functions.

TABLE G.1
Sequenced Agenda for Emergency Managers and Others Who Support EM Functions

Time Horizon	Agenda
Next 12 months	• Pursue opportunities to build IT knowledge and skill, including AI specifically. • Assess your organization's risk tolerance, including go and no-go indicators to watch for. • Learn what AI-enabled products are available to support your work. • Assess your organization's readiness for AI-enabled products using Appendix H as a resource. • Track AI governance developments inside your organization and at higher levels of government and consider what they mean for using AI-enabled products to support EM. • Share implementation choices and experiences (e.g., for AI products being tried, configured, or rejected) through professional associations, conferences, and informal networks. • Get involved in and support peer networks through which diffusion can spread. • Document your organization's data and connectivity conditions before evaluating products.
Next 1 to 3 years	• Build AI considerations into tactical, operational, and strategic planning efforts. • Pilot AI-enabled products to support defined tasks. • Interview technology companies with products that look useful using Appendix I to guide those discussions.
Ongoing	• Seek opportunities to collaborate with technology companies that are interested in user-centered design. • Evaluate AI-enabled product implementations in your office or organization, identify what is and is not working, and feed that learning back into use. • Encourage professional development partners to include AI training, education, and peer learning in curriculum and agendas.

An Agenda for Technology Companies

The landscape documented in Chapter 2 is substantial. Technology companies have built that landscape, and they continue to expand it. They are not the organizations that are adopting AI-enabled products to support EM, but their decisions materially shape whether the products they offer can be broadly adopted and reach the point of diffusion (Chapter 4).

Three patterns from Chapters 2 and 3 are relevant here. First, product information is often fragmented and inconsistent and not publicly available. Second, many technology companies implicitly assume that a customer has the procurement bandwidth and IT capacity to stand up their products. Technology company structures are often imperfectly aligned with EM operating realities, and many companies do not offer 24/7 technical assistance. Technology companies control the levers that shape these patterns, and Table G.2 provides a sequenced action agenda across three timelines for using these levers.

However, not all companies are positioned to act on these levers in the same way. Big tech and hyperscalers operate at a scale and from a resource base that make it possible to take certain actions that smaller technology companies cannot. Depending on what happens with federal funding for EM at the SLTT level, big tech and hyperscalers may be uniquely positioned to support AI adoption and diffusion in EM at a scale no other entities can match (see Table G.3 for a sequenced action agenda across three timelines and the "High-Impact Initiatives to Accelerate Adoption and Diffusion at Scale" section in Chapter 5).

TABLE G.2

Sequenced Agenda for Technology Companies

Time Horizon	Agenda
Next 12 months	• Make understanding the EM landscape, its needs, and its limitations a priority effort. • Begin collaboration among companies offering purpose-built EM products to develop product communication standards. • Make pricing visible and predictable. • Make dependencies, integration and data requirements, and the total cost of ownership visible at the point of marketing. • Make security and governance documentation easy to find.
Next 1 to 3 years	• Work with standards bodies and industry associations to move toward standardized disclosure on the dimensions buyers need to evaluate most. • Develop new products that are consistent with the design criteria in this report, scaled to company resources. • Develop products that address the gaps documented in the landscape. • Configure existing general-purpose technologies for EM use, either as embedded features in EM products or as stand-alone offerings (for example, adding AI-enabled translation capability to a warning system). • Invest in 24/7 support, particularly for products meant to function during response.
Ongoing	• Participate in groups that connect technology companies with emergency managers and others who support EM functions. • Provide standardized product information in key areas. • Reduce the integration and ongoing data burden during product design. • Ensure that response-supporting products work with intermittent or no internet access. • Address risks documented in the landscape. • Support AI governance and pursue AI governance certification.

TABLE G.3
Sequenced Agenda for Big Tech and Hyperscalers

Time Horizon	Agenda
Next 12 months	• In cooperation with other stakeholder groups, – spearhead the development of a standing cross-sector alliance – launch an EM AI readiness and adoption accelerator – fund adoption-relevant research gaps – fund or otherwise support an education and training program to increase awareness and knowledge.
Next 1 to 3 years	• Develop AI-enabled products that are purpose-built for EM and that meet EM-centered design criteria (e.g., light or no integration burden, no ongoing user data-provision requirement, mature, low risk, cyber certified, cheap or free). • Examine what more-general purpose products you could configure for EM task areas when purpose building is not the path. • Do demonstration projects to illustrate to technology companies innovative ways to offer integrated solutions that reduce the stacking burden to support EM functions. • Offer products at no or low cost to capacity-constrained EM buyers.
Ongoing	• Engage with FEMA and EM professional associations on national AI strategies. • Fund the creation and sustainment of a cross-sector alliance to support focused attention to the adoption and diffusion of AI-enabled products and propel and sustain progress toward these goals. • Fund efforts to improve AI and technical literacy in the EM-serving workforce. • Fund research efforts to support AI adoption and diffusion.

An Agenda for Policymakers

In a government context, elected officials determine whether EM offices and the organizations that support EM functions are adequately funded and fully staffed for the work expected of them, and they approve policy (e.g., AI governance). Through their priority setting and decisionmaking, elected officials are among the policymakers who shape the extent to which AI-enabled products will be accessed and used and how quickly that happens. Public-sector executives who decide what their organizations procure and deploy are also in this group, including IT department leaders who manage the systems on which AI-enabled products depend and procurement leaders who structure how products are acquired. Outside the government context, executive leadership and those who are in charge of IT and finance policy play these roles. Regardless of context, policymaker decisions sit at the center of the conditions documented in Chapter 4: They can remove the barriers that emergency managers and others who support EM functions cannot. Table G.4 lays out a sequenced action agenda across three time horizons for policymakers and public-sector leaders.

TABLE G.4
Sequenced Agenda for Policymakers

Time Horizon	Agenda
Next 12 months	• Fund technology specifically within EM offices, including the procurement, implementation, and sustainment of AI-enabled products and the IT and staff capacity to use them. • When funding technology in organizations that support EM functions, require consideration of EM use cases in the process. • When the surrounding entity procures AI-enabled products for staff use, extend access to the EM office it houses. • Proactively share purchasing pathways, including cooperative contracts, with EM offices and the organizations that support EM functions.
Next 1 to 3 years	• Build AI-specific procurement infrastructure, including AI procurement guidance, model contract terms, and review pathways. • Develop AI governance that accounts for EM operational realities, including 24/7 operations and data sharing with response partners. • Where state-to-local AI provisioning is possible, establish or expand it. • Expand the services that state IT organizations extend to local governments to include support for AI-enabled product evaluation and adoption, not only cybersecurity.
Ongoing	• Review current and developing AI policy for unintended consequences for EM use. • When governing AI use and AI infrastructure, account for EM dependencies on both. • Require coordination across IT, procurement, EM, and the surrounding entity.

An Agenda for Researchers

Researchers play a significant role in what information is available to other stakeholder groups to know and act on. Chapter 4 documented uneven evidence on AI adoption in EM offices, in the organizations that support EM functions, and across the policy and procurement systems that govern AI use, with the strongest available evidence often coming from outside EM and applied by analogy. The research community has significant influence over the production and direction of that evidence: what questions get asked, which populations are studied, how findings are framed for use, and whether the resulting work reaches stakeholders. Table G.5 lays out a sequenced action agenda across three time horizons for researchers.

TABLE G.5
Sequenced Agenda for Researchers

Time Horizon	Agenda
Next 12 months	• Make relevant research findings accessible to the practitioner audience. • Identify and publish a research agenda to support the adoption and diffusion of AI-enabled products. • Document how AI-enabled products are being used now by emergency managers and others that support EM functions.
Next 1 to 3 years	• Provide answers to the following key research questions: – Why have long-standing deployed products not yet reached full adoption? – How are AI-enabled products diffused across EM-related fields? – What outcomes result from the use of AI-enabled products in EM functions where implemented? – Can LLMs effectively support EM tasks—and under what conditions—based on established effectiveness criteria in the hazards and disaster research? – Could national AI governance, procurement, and IT policies create systematic barriers to adoption, and what policy changes could address those barriers?
Ongoing	• Participate in collaborative groups (e.g., cross sector alliance) to develop an AI for EM research agenda that is grounded in real-world adoption and diffusion needs. • Clearly define the study population, recognizing that EM is not a single system, and account for differences across EM offices and the broader fields that support EM operations.

APPENDIX H

Criteria Practitioners Can Use to Evaluate Their Organization's Adoption Readiness Based on Product Characteristics

In this appendix, we provide a list of questions to help practitioners deliberate about their own organization's readiness to procure and adopt AI-enabled products given the considerations found in this research. Not all criteria apply equally. Each organization should weigh each domain based on its unique context, constraints, and requirements. Use this assessment to identify go and no-go thresholds for your specific situation. The questions in Figure H.1 were derived from the barriers and challenges identified in this research and grounded in practitioner expertise, technology adoption frameworks, and organizational readiness principles.

FIGURE H.1

Questions About Readiness to Procure and Adopt AI-Enabled Products

01 APPLICABILITY READINESS

Are we solving the right problem with the right product?

CONSIDERATION	READY	PARTIALLY READY	NOT READY
Do we have clearly defined workflows and decision points that this product would support?	○	○	○
Are our targeted EM functions for this product (e.g., situational awareness, recovery coordination) consistently performed or still ad hoc?	○	○	○
Do we have clarity on what *better* looks like for this task or functional area (e.g., speed, accuracy, coordination, outcomes) in order to evaluate whether a product would be worth pursuing?	○	○	○
Are we aligned on the problem(s) we are trying to solve? And/or the opportunities we are trying to advance?	○	○	○

02 ADOPTABILITY READINESS

Can our workforce take this on?

CONSIDERATION	READY	PARTIALLY READY	NOT READY
Do we have staff with the time and bandwidth to adopt and use a product during steady state and/or incidents?	○	○	○
Do we have cultural, philosophical, or emotional barriers to overcome to use AI in general as an emerging technology?	○	○	○
Do staff have the baseline technical proficiency required?	○	○	○
Is there a willingness to change workflows or a strong attachment to current practices?	○	○	○
Do we have internal champions to drive adoption and troubleshoot early issues?	○	○	○
Can we sustain training, onboarding, and turnover for a new technology product?	○	○	○

03 ADAPTABILITY READINESS

Are our systems and processes flexible enough?

CONSIDERATION	READY	PARTIALLY READY	NOT READY
Are our workflows and SOPs flexible enough to incorporate new products?	○	○	○
Can we tolerate iterative refinement (e.g., testing, adjusting thresholds, tuning outputs)?	○	○	○
Are we able to scale operations up or down in ways that an AI product might expect?	○	○	○

04 INTEGRATION AND DEPENDENCY READINESS

Do we have the technical foundation in place?

CONSIDERATION	READY	PARTIALLY READY	NOT READY
Do we have the underlying systems the product depends on (e.g., GIS, data platforms, alerting systems)?	○	○	○
Is our data accessible, structured, and governed well enough to support the product?	○	○	○
Do we understand our data sources, quality, and gaps?	○	○	○
Can we support integrations (e.g., APIs, data feeds) or will that require external support?	○	○	○
Are we prepared for stacking requirements (i.e., this product being one component in a broader ecosystem)?	○	○	○

05 RELIABILITY AND OPERATIONAL READINESS

Can we use this safely during real incidents?

CONSIDERATION	READY	PARTIALLY READY	NOT READY
Do we have defined fallback procedures if the product is unavailable during an incident?	○	○	○
Can we monitor and validate outputs before acting on them?	○	○	○
Do we have an operational culture that supports cross-checking and not over-relying on automation?	○	○	○
Are we able to incorporate this product into exercises, drills, and AARs?	○	○	○

06 GOVERNANCE, LEGAL, AND CYBER READINESS

Are the right policies and stakeholders in place?

CONSIDERATION	READY	PARTIALLY READY	NOT READY
Do we have policies for data use, sharing, and retention that align with this product?	○	○	○
Is there clarity on decision authority when using automated or AI-supported outputs?	○	○	○
Do we have cybersecurity practices that are sufficient to manage a new connected system?	○	○	○
Are legal, procurement, and IT stakeholders engaged early enough to support adoption?	○	○	○

07 PROCUREMENT AND RESOURCE READINESS

Can we afford and sustain this?

CONSIDERATION	READY	PARTIALLY READY	NOT READY
Do we have a viable funding model that matches the product (e.g., subscription versus capital purchase)?	○	○	○
Can we sustain costs over multiple years, including maintenance and upgrades?	○	○	○
Are we able to navigate procurement pathways—contracts, approvals, and timelines?	○	○	○
Do we understand the implications of vendor lock-in or long-term dependencies?	○	○	○

08 MATURITY AND LEARNING ORIENTATION

Are we ready to learn and improve as we go?

CONSIDERATION	READY	PARTIALLY READY	NOT READY
Have we successfully adopted similar products in the past? What were the lessons?	○	○	○
Do we have a culture of continuous improvement and a willingness to iterate?	○	○	○
Are we prepared to measure outcomes and adjust our use based on actual performance?	○	○	○
Do we have relationships with peer jurisdictions to share lessons and practices?	○	○	○

APPENDIX I

Questions for Product Technology Companies: A Checklist for Practitioners

In this appendix, we provide practitioners, their procurement offices, and/or their IT sections with questions they can ask technology companies when evaluating a potential product for procurement. Depending on several organizational factors, some questions will be more important than others (e.g., go or no-go criteria could depend on the unique constraints and requirements of each organization). The organization may also consider creating a customized scorecard that gives a weight to each of the categories below and identifies any clear go or no-go criteria. Additionally, the questions presented in Figure I.1 are intended to serve as a menu of inquiries that practitioner teams can adapt based on their priorities and specific circumstances.

FIGURE I.1

Questions About the Applicability of AI-Enabled Products

Mark any question as a go or no-go criterion

01 APPLICABILITY

Are we solving the right problem with the right product?

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ What specific EM functions or task areas does this product support (e.g., situational awareness, incident coordination, damage assessments)?	
☐ Is the product hazard specific, agnostic, or for selected hazards?	
☐ Where does this capability fit within a steady state, an incident, and/or a recovery workflow?	
☐ What decisions does this product inform or improve?	
☐ What measurable outcomes have been demonstrated in EM or closely related domains?	
☐ What category of AI underlies this product (e.g., LLM, statistical or predictive ML, physics-based simulation, CV, or a combination), and which component is responsible for which function?	

02 ADOPTABILITY

What requirements will this product impose on my agency?

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ What staffing, training, and onboarding are required? What is available from your company?	
☐ How intuitive is the interface for nontechnical users under time pressure?	
☐ What change management challenges have prior users encountered?	
☐ How long does it typically take from procurement to operational use?	
☐ What assumptions does the product make about the operating environment (e.g., data availability, staffing, tempo)?	
☐ What internal roles (e.g., IT, GIS, analysts) are required to sustain it?	
☐ What governmental policies or guidance have affected procurement or implementation among other users?	

03 ADAPTABILITY

How well does this product interface with my unique environment?

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ How configurable is the system to local hazards, jurisdictions, and operational structures?	
☐ Can users modify workflows, thresholds, or outputs without vendor intervention?	
☐ How well does the system operate across preparedness, response, and recovery?	
☐ How does the product perform under degraded, incomplete, or conflicting data conditions?	
☐ Can the system scale with incident size and complexity?	

04 RELIABILITY AND SUPPORT

What should I expect in terms of reliability and vendor support?

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ Who developed the underlying model or models? If any are licensed or accessed through a third party, which providers are involved? When was the model last updated, and how are version changes communicated to customers?	
☐ How is output correctness evaluated before release and on an ongoing basis? Specifically, what test sets are used, what is the measured performance against them, who designs and reviews the tests, and what failure modes have been documented?	
☐ What is the system's performance under incident conditions (e.g., latency, uptime, load)?	
☐ What is the error profile (e.g., false positives or negatives), and how is it measured?	
☐ How transparent and explainable are outputs and recommendations?	
☐ What are the fallback procedures if the system fails or produces bad outputs?	
☐ What level of operational support is available (e.g., 24/7, surge support, business hours)?	

05 PRIVACY, LEGAL, AND CYBERSECURITY

What potential risks am I willing to accept from the vendor and/or product?

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ What data is collected, stored, and transmitted, and where is it hosted?	
☐ Is jurisdictional data used to train or improve the vendor's models?	
☐ What cybersecurity standards and certifications are met?	
☐ How are access controls, audit logs, and data protections managed?	
☐ What legal or liability considerations have arisen in past uses?	
☐ Are the vendor's commitments on training data use, retention periods, and data residency contractually binding on any upstream model providers, hosting partners, or subprocessors whose infrastructure handles customer content? Where in the contract is this established?	

05 PRIVACY, LEGAL, AND CYBERSECURITY

What potential risks am I willing to accept from the vendor and/or product?

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ Does the product include any current or planned capability that allows customer content, derived outputs, embeddings, or usage data to be visible, retrievable, or otherwise accessible to other customers, whether by default, by opt in, or in anonymized form? If yes, describe the mechanism, the scope of what is shared, and the customer's controls over it.	
☐ If a customer contributes to a document or any shared or pooled feature and later withdraws that contribution, what specifically is removed, on what timeline, and what happens to the outputs other customers have already generated (e.g., saved summaries, exported reports, prior chat responses, downstream analyses)?	

06 PROCUREMENT

What are the true expected costs, purchase models, and dependencies

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ What procurement models are available (e.g., one-time license, subscription or SaaS, usage based, enterprise agreement)?	
☐ What is the total cost of ownership over three to five years (e.g., integration, data, sustainment, upgrades)?	
☐ What dependencies exist on other products, platforms, or data providers?	
☐ Does effective use require stacking with additional products (e.g., GIS, data platforms, alerting systems)?	
☐ Are there proprietary components that create vendor lock-in or limit interoperability?	
☐ What integration work is required with existing systems (e.g., WebEOC, CAD, GIS, alerting)?	
☐ What are the data ownership, portability, and exit provisions if the contract ends?	
☐ How are updates, versioning, and feature changes handled over time?	

06 PROCUREMENT

What are the true expected costs, purchase models, and dependencies

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ What procurement pathway has been used by other public-sector customers?	
☐ For each major commitment made during the sales process (e.g., training data use, retention, security posture, accuracy, support availability, integration scope), where exactly is that commitment documented in terms of service, data processing, or master services agreements? Identify the specific clauses. This is the demonstration-to-contract reconciliation move. It puts the burden on the vendor to point to written language.	

06 PROCUREMENT

What are the true expected costs, purchase models, and dependencies

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ How are middle-term changes to the underlying AI model, upstream providers, or core product behavior governed contractually? Specifically, what notice is required, what customer rights are triggered (e.g., review, opt out, termination), and what happens to data residency or processing terms if the underlying provider changes? This catches the *silent swap* risk. A vendor that replaces their LLM provider mid-contract can change behavior, accuracy, the pricing structure, and data residency without you noticing.	
☐ What is the vendor's liability framework, specifically for AI-generated errors that cause operational consequences (e.g., a flawed recommendation acted on during an incident)? What is the liability cap, what carve-outs apply, and what contractual language preserves practitioner judgment as the decision of record?	

07 MATURITY AND REFERENCES

What independent evidence exists that this product delivers as promised?

QUESTIONS TO ASK THE VENDOR	NOTES AND VENDOR RESPONSES
☐ What real-world deployments exist in comparable jurisdictions or incidents?	
☐ Can you provide documented use cases with outcomes (e.g., time saved, decisions improved)?	
☐ How long has the product been in operational use?	
☐ What is the vendor's track record for updates, maintenance, and support?	
☐ Who are reference customers we can speak with directly?	
☐ What published policies, frameworks, or external standards govern the vendor's responsible AI practices (e.g., NIST's AI Risk Management Framework, ISO/IEC 42001, sector-specific federal guidance)? Provide an example of a feature, capability, or use case the vendor has chosen not to support and the reasoning behind that decision.	

Abbreviations

AAR	after-action report
ADMS	advanced distribution management system
AG	attorney general
AI	artificial intelligence
AIDE	AI for Disasters and Emergencies
API	application programming interface
BEOC	business emergency operations center
CAD	computer-aided dispatch
CRM	customer relationship management
CV	computer vision
EM	emergency management
EMS	emergency medical services
EOC	emergency operations center
ETA	estimated time of arrival
FCC	Federal Communications Commission
FedRAMP	Federal Risk and Authorization Management Program
FEMA	Federal Emergency Management Agency
FTC	Federal Trade Commission
GIS	geographic information system
GPS	Global Positioning System
GPT	generative pretrained transformer
HSRD	Homeland Security Research Division
ICS	incident command system
IEC	International Electrotechnical Commission
IoT	Internet of Things
ISO	International Organization for Standardization
IT	information technology
LLM	large language model
ML	machine learning
NGO	nongovernmental organization
NIMS	National Incident Management System
NIST	National Institute of Standards and Technology
NLP	natural language processing
on-prem	on-premises
OSINT	open-source intelligence
P3	public-private partnership
PII	personally identifiable information
Q&A	question and answer

SaaS	software as a service
SAM	system for award management
SAR	synthetic aperture radar
SAVER	System Assessment and Validation for Emergency Responders
SBA	Small Business Administration
SEC	Securities and Exchange Commission
SLTT	state, local, tribal, and territorial
SOC	System and Organization Controls
SOP	standard operating procedure

References

Achiam, Josh, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al., "GPT-4 Technical Report," arXiv, arXiv:2303.08774, March 15, 2023.

Administrative Office of the U.S. Courts, *Judicial Business of the United States Courts: 2023 Annual Report*, U.S. Courts, 2024.

AIDE—*See* AI for Disasters and Emergencies.

AI for Disasters and Emergencies, homepage, undated. As of April 21, 2026: https://www.aideinitiative.org/

Aikins, Steven Kwamena, and Tamara Dimitrijevska-Markoski, "Determinants of Government AI Adoption for e-Government," *AI & Society*, 2026.

Anil, Rohan, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, et al., "Gemini: A Family of Highly Capable Multimodal Models," arXiv, arXiv:2312.11805, December 19, 2023.

Anthropic, "Claude Code by Anthropic," webpage, undated. As of May 11, 2026: https://www.anthropic.com/product/claude-code.

Anthropic, "Introducing the Next Generation of Claude," webpage, March 4, 2024. As of April 29, 2026: https://www.anthropic.com/news/claude-3-family

Apple Machine Learning Research, "Hey Siri: An On-Device DNN-Powered Voice Trigger for Apple's Personal Assistant," October 1, 2017.

Aranda, Luis, and Kasumi Sugimoto, "The Agentic AI Landscape and Its Conceptual Foundations," Organisation for Economic Co-operation and Development, OECD Artificial Intelligence Papers No. 56, February 2026.

Barr, Jonathan, Alexander Hagen, Grant Tietje, Nicholas Betzsold, Alex Greer, and Eric Best, *Emergency Management of Tomorrow Research: Artificial Intelligence Landscape Assessment*, Science and Technology Directorate, U.S. Department of Homeland Security, and Pacific Northwest National Laboratory, June 2024.

Basu, Ragini, and Trina Sheets, *Deloitte–NEMA National Risk Study 2025: Changing Landscapes in State Emergency Management*, Deloitte Center for Government Insights, Deloitte, 2025.

Bishop, Christopher M., *Pattern Recognition and Machine Learning*, Springer, 2006.

Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al., "On the Opportunities and Risks of Foundation Models," arXiv, arXiv:2108.07258, 2021.

Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al., "Language Models Are Few-Shot Learners," in Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, eds., *Advances in Neural Information Processing Systems*, Vol. 33, 2020.

Center for Digital Government, *2024 Generative AI Survey*, e.Republic, May 2024.

Chen, Tzuhao, Mila Gascó-Hernandez, and Marc Esteve, "The Adoption and Implementation of Artificial Intelligence Chatbots in Public Organizations: Evidence from U.S. State Governments," *American Review of Public Administration*, Vol. 54, No. 3, April 2024.

Coggburn, Jerrell D., "Exploring Differences in the American States' Procurement Practices," *Journal of Public Procurement*, Vol. 3, No. 1, 2003.

Currie, Chris, *Disaster Assistance High-Risk Series: State and Local Response Capabilities*, U.S. Government Accountability Office, GAO-26-108599, December 18, 2025.

Damanpour, Fariborz, and Marguerite Schneider, "Phases of the Adoption of Innovation in Organizations: Effects of Environment, Organization and Top Managers," *British Journal of Management*, Vol. 17, No. 3, September 2006.

David, Anne, Tan Yigitcanlar, Rita Yi Man Li, Juan M. Corchado, Pauline Hope Cheong, Karen Mossberger, and Rashid Mehmood, "Understanding Local Government Digital Technology Adoption Strategies: A PRISMA Review," *Sustainability*, Vol. 15, No. 12, 2023.

DiMaggio, Paul J., and Walter W. Powell, "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields," *American Sociological Review*, Vol. 48, No. 2, April 1983.

Di Troia, Sarah, Vanessa Parli, Juan N. Pava, Haifa Badi Uz Zaman, and Kelly Fitzsimmons, "Inspiring Action: Identifying the Social Sector AI Opportunity Gap," Institute for Human-Centered Artificial Intelligence, Stanford University, and Project Evident, Working Paper, February 2024.

Drabek, Thomas E., "Microcomputer Usage in Disaster Preparedness and Response," *Industrial Crisis Quarterly*, Vol. 5, No. 2, June 1991.

Emergency Management Accreditation Program, *Emergency Management Standard*, May 2022.

Enrich Layer, homepage, undated. As of June 12, 2026:
https://enrichlayer.com/

European Data Protection Board, *Annual Report 2023: Safeguarding Individuals' Digital Rights*, 2024.

Federal Emergency Management Agency, "FEMA Geospatial Resource Center," webpage, undated. As of April 29, 2026:
https://gis-fema.hub.arcgis.com/

Federal Emergency Management Agency, "Principles of Emergency Management," 2007.

Federal Emergency Management Agency, *National Preparedness Goal*, 2nd ed., U.S. Department of Homeland Security, September 2015.

Federal Emergency Management Agency, *National Disaster Recovery Framework*, 2nd ed., U.S. Department of Homeland Security, June 2016.

Federal Emergency Management Agency, *National Response Framework*, 4th ed., U.S. Department of Homeland Security, October 28, 2019.

Federal Emergency Management Agency, *Strategic Foresight 2050: Final Report*, November 2024.

Federal Emergency Management Agency, "Resilience Analysis and Planning Tool (RAPT)," webpage, last updated April 8, 2026. As of April 29, 2026:
https://www.fema.gov/emergency-managers/practitioners/resilience-analysis-and-planning-tool

Feigenbaum, Edward A., "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Vol. 2, 1977.

Feigenbaum, Edward A., and Pamela McCorduck, *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, 1st ed., Addison-Wesley, 1983.

FEMA—*See* Federal Emergency Management Agency.

Fichman, Robert G., and Chris F. Kemerer, "The Assimilation of Software Process Innovations: An Organizational Learning Perspective," *Management Science*, Vol. 43, No. 10, October 1997.

Forrest, Page, and Jad Maayah, "Uncertainty Surrounding Federal Disaster Funding Looms over State Budgets," Pew Charitable Trusts, January 15, 2026.

Gallivan, Michael J., "Organizational Adoption and Assimilation of Complex Technological Innovations: Development and Application of a New Framework," *Data Base for Advances in Information Systems*, Vol. 32, No. 3, Summer 2001.

Gao, Lu, and Luis Felipe Luna-Reyes, "Implications of Adopting AI-Based Applications in Local Governments: A Systematic Literature Review from a Public Value Perspective," in Younhee Kim and Michael J. Ahn, eds., *The Art of Digital Governance: Navigating Platforms and AI Revolution*, Springer, 2025.

Glasscock, Amy, Emily Lane, Eric Sweden, Alex Whitaker, and Kalea Young-Gibson, *The 2025 State CIO Survey: Leading Change Through Uncertain Times*, National Association of State Chief Information Officers, 2025.

Goodfellow, Ian J., Jean Pouget-Abadiey, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozairz, Aaron Courville, and Yoshua Bengio, "Generative Adversarial Nets," in Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger, eds., *Advances in Neural Information Processing Systems*, Vol. 27, 2014.

Greenhalgh, Trisha, Glenn Robert, Fraser Macfarlane, Paul Bate, and Olivia Kyriakidou, "Diffusion of Innovations in Service Organizations: Systematic Review and Recommendations," *Milbank Quarterly*, Vol. 82, No. 4, December 2004.

Hendrycks, Dan, Mantas Mazeika, and Thomas Woodside, "An Overview of Catastrophic AI Risks," arXiv, arXiv:2306.12001v6, October 9, 2023.

Hinkley, Sara, *Technology in the Public Sector and the Future of Government Work*, UC Berkeley Labor Center, University of California, Berkeley, January 2023.

Ho, Jonathan, Ajay Jain, and Pieter Abbeel, "Denoising Diffusion Probabilistic Models," in Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, *Advances in Neural Information Processing Systems*, Vol. 33, 2020.

Hofstadter, Douglas R., *Gödel, Escher, Bach: An Eternal Golden Braid*, Basic Books, 1979.

International Organization for Standardization and International Electrotechnical Commission, *Information Technology—Artificial Intelligence—Management System*, ISO/IEC 42001:2023, 1st ed., December 2023.

ISO and IEC—*See* International Organization for Standardization and International Electrotechnical Commission.

Jackson, Peter, *Introduction to Expert Systems*, 3rd ed., Addison-Wesley, 1998.

Jennings, Eliot A., Sudha Arlikatti, Simon A. Andrew, and KyungWoo Kim, "Adoption of Information and Communication Technologies (ICTs) by Local Emergency Management Agencies in the United States," *International Review of Public Administration*, Vol. 22, No. 2, 2017.

Jensen, Jessica, Sarah Bundy, Brian Thomas, and Mariama Yakubu, "The County Emergency Manager's Role in Disaster Recovery," *International Journal of Mass Emergencies and Disasters*, Vol. 32, No. 1, March 2014.

Ji, Ziwei, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung, "Survey of Hallucination in Natural Language Generation," *ACM Computing Surveys*, Vol. 55, No. 12, March 2023.

Johnson, Nari, Elise Silva, Harrison Leon, Motahhare Eslami, Beth Schwanke, Ravit Dotan, and Hoda Heidari, "Legacy Procurement Practices Shape How U.S. Cities Govern AI: Understanding Government Employees' Practices, Challenges, and Needs," *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, Association for Computing Machinery, 2025.

Labadie, John R., "Problems in Local Emergency Management," *Environmental Management*, Vol. 8, No. 6, November 1984.

LaLone, Nicolas, Phoebe O. Toups Dugas, and Bryan Semaan, "The Technology Crisis in US-Based Emergency Management: Toward a Well-Connected Future," in Tung X. Bui, ed., *Proceedings of the 56th Hawaii International Conference on System Sciences*, University of Hawaii, 2023.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton, "Deep Learning," *Nature*, Vol. 521, 2015.

Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-Tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, *Advances in Neural Information Processing Systems*, Vol. 33, 2020.

Lucci, Stephen, Sarhan M. Musa, and Danny Kopec, *Artificial Intelligence in the 21st Century: A Living Introduction*, 3rd ed., Mercury Learning and Information, 2022.

Markle Foundation, "About Markle," webpage, undated. As of April 21, 2026: https://www.markle.org/about-markle/

Markle Foundation, "The Markle Foundation Launches AI Initiative Focused on Emergency Preparedness, Response, and Recovery," December 10, 2025.

Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njenga Kariuki, Emily Capstick, Anka Reuel, Erik Brynjolfsson, John Etchemendy, et al., *Artificial Intelligence Index Report 2025*, AI Index Steering Committee, Institute for Human-Centered Artificial Intelligence, Stanford University, April 2025.

McCarthy, John, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," *AI Magazine*, Vol. 27, No. 4, [1955] 2006.

McCorduck, Pamela, *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*, 2nd ed., A K Peters, 2004.

McKinney, Wes, "Data Structures for Statistical Computing in Python," in Stéfan van der Walt and Jarrod Millman, eds., *Proceedings of the 9th Python in Science Conference*, 2010.

McTarnaghan, Sara, Lucy Dadayan, and Andrew Rumbach, "The Trump Administration Wants to Shrink the Federal Government's Role in Disaster Management. Do States Have the Fiscal Capacity to Weather the Storm?" Urban Institute, July 15, 2025.

Meyer, John W., and Brian Rowan, "Institutionalized Organizations: Formal Structure as Myth and Ceremony," *American Journal of Sociology*, Vol. 83, No. 2, September 1977.

Mitchell, Tom M., *Machine Learning*, 1st ed., McGraw-Hill, 1997.

National Association of State Procurement Officials, *2024 Survey of State Procurement Practices Report*, 2024.

National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, U.S. Department of Commerce, January 2023.

National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1, U.S. Department of Commerce, July 2024.

National Preparedness Analytics Center, *Emergency Management Organizational Structures, Staffing, and Capacity Study: State, Local, and Territory Findings Report*, Argonne National Laboratory, July 2025.

National Research Council, *Improving Disaster Management: The Role of IT in Mitigation, Preparedness, Response, and Recovery*, National Academies Press, 2007.

Natural Hazards Center, "State of Emergency—Resource Commitments and Compassion in an Age of Extremes," video, July 13, 2023. As of June 1, 2026:
https://youtu.be/cQT_MuWiIzU

Nilsson, Nils J., *Artificial Intelligence: A New Synthesis*, Morgan Kaufmann, 1998.

NIST—*See* National Institute of Standards and Technology.

Obama, Barack, Presidential Policy Directive 8: National Preparedness, March 30, 2011.

OpenAI, "Introducing ChatGPT," webpage, November 30, 2022. As of April 29, 2026:
https://openai.com/blog/chatgpt

OpenAI, "Introducing Codex," webpage, May 16, 2025. As of May 11, 2026:
https://openai.com/index/introducing-codex/

Petak, William J., "Emergency Management: A Challenge for Public Administration," *Public Administration Review*, Vol. 45, January 1985.

Pothireddy, Sudheekar Reddy, "From Data to Impact: How AI and Microsoft Enhance Efficiency in Nonprofit Organizations," *International Journal for Multidisciplinary Research*, Vol. 7, No. 1, January–February 2025.

Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever, "Language Models Are Unsupervised Multitask Learners," OpenAI, 2019.

Ratan, Ishana, Alison E. Post, Tanu Kumar, and Mridang Sheth, "Which Local Governments Adopt New Technology First? Agency Size and Organizational Champions for Open Transit Data," *Urban Affairs Review*, Vol. 62, No. 62, March 2026.

Reddick, Christopher, "Information Technology and Emergency Management: Preparedness and Planning in US States," *Disasters*, Vol. 35, No. 1, January 2011.

Rogers, Everett M., *Diffusion of Innovations*, 5th ed., Free Press, 2003.

Rotter, Janne, and William Bailkoski, "AI Adoption in NGOs: A Systematic Literature Review," arXiv, arXiv:2510.15509v1, October 17, 2025.

Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams, "Learning Representations by Back-Propagating Errors," *Nature*, Vol. 323, 1986.

Russell, Stuart, and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed., Pearson, 2022.

Sajadieh, Sha, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Lapo Santarlasci, Juan Pava, Nestor Maslej, Russ Altman, Erik Brynjolfsson, et al., *Artificial Intelligence Index Report 2026*, AI Index Steering Committee, Institute for Human-Centered Artificial Intelligence, Stanford University, April 2026.

Salesforce, "Agentforce for Energy & Utilities," webpage, undated-a. As of July 13, 2026: https://www.salesforce.com/energy-utilities/artificial-intelligence/

Salesforce, "Emergency Response Management," webpage, undated-b. As of July 13, 2026: https://www.salesforce.com/government/guided-tours/emergency-response/

Samuelson, Pamela, "Generative AI Meets Copyright," *Science*, Vol. 381, No. 6654, 2023.

Sarikaya, Ruhi, "The Technology Behind Personal Digital Assistants: An Overview of the System Architecture and Key Components," *IEEE Signal Processing Magazine*, Vol. 34, No. 1, January 2017.

Science and Technology Directorate, U.S. Department of Homeland Security, "System Assessment and Validation for Emergency Responders (SAVER)," webpage, last updated June 10, 2026. As of June 12, 2026: https://www.dhs.gov/science-and-technology/saver

ServiceNow, *Impact AI: Nonprofit Digital Transformation Report 2024*, 2024.

Shavit, Yonadav, Sandhini Agarwal, Miles Brundage, Steven Adler, Cullen O'Keefe, Rosie Campbell, Teddy Lee, Pamela Mishkin, Tyna Eloundou, Alan Hickey, et al., "Practices for Governing Agentic AI Systems," OpenAI, 2023.

Shi, Wanzhu, and Lauren Azevedo, "Determinants of AI Adoption in Nonprofit Organizations," *Nonprofit and Voluntary Sector Quarterly*, 2026.

Smith Arrillaga, Elisha, Seara Grundhoefer, and Christina Im, *AI with Purpose: How Foundations and Nonprofits Are Thinking About and Using Artificial Intelligence*, Center for Effective Philanthropy, 2025.

Sohl-Dickstein, Jascha, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli, "Deep Unsupervised Learning Using Nonequilibrium Thermodynamics," in Francis Bach and David Blei, eds., *Proceedings of the 32nd International Conference on Machine Learning*, Vol. 37, 2015.

Tolbert, Pamela S., and Lynne G. Zucker, "Institutional Sources of Change in the Formal Structure of Organizations: The Diffusion of Civil Service Reform, 1880–1935," *Administrative Science Quarterly*, Vol. 28, No. 1, March 1983.

Turing, A. M., *Proposals for Development in the Mathematics Division of an Automatic Computing Engine*, National Physical Laboratory, E.882, 1945.

U.S. Department of Homeland Security, *Department of Homeland Security Artificial Intelligence Roadmap 2024*, 2024.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin, "Attention Is All You Need," in Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, eds., *Advances in Neural Information Processing Systems*, Vol. 30, 2017.

Virtuous and Fundraising.AI, *The 2026 Nonprofit AI Adoption Report: A Benchmark Study of 346 Organizations*, Virtuous, February 2026.

Weizenbaum, Joseph, "ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine," *Communications of the ACM*, Vol. 9, No. 1, 1966.

Wickham, Hadley, Mara Averick, Jennifer Bryan, Winston Chang, Lucy D'Agostino McGowan, Romain François, Garrett Grolemund, Alex Hayes, Lionel Henry, Jim Hester, et al., "Welcome to the Tidyverse," *Journal of Open Source Software*, Vol. 4, No. 43, 2019.

World Economic Forum, *The Global Risks Report 2023*, 18th ed., January 2023.

Xiao Liang, "Image-Based Post-Disaster Inspection of Reinforced Concrete Bridge Systems Using Deep Learning with Bayesian Optimization," *Computer-Aided Civil and Infrastructure Engineering*, Vol. 34, No. 5, May 2019.

Young, Gerald, *Artificial Intelligence in the Workforce: A Survey of State and Local Employees*, MissionSquare Research Institute, June 2025.

Young, Shalanda D., "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," memorandum for the heads of executive departments and agencies, Office of Management and Budget, M-24-10, March 28, 2024.

Zhu, Kevin, Kenneth L. Kraemer, and Sean Xu, "The Process of Innovation Assimilation by Firms in Different Countries: A Technology Diffusion Perspective on E-Business," *Management Science*, Vol. 52, No. 10, October 2006.