WHITE PAPER

The Augmented Planner

Scaling SLTT Emergency Management Capability Through Human-Centered AI

Sentinel Resilience Partners | Arlington, Virginia | sentinelresiliencepartners.com | June 2026

Executive Summary

State, Local, Tribal, and Territorial (SLTT) emergency management agencies are being asked to do more with less than at any point in the profession's modern history. In a 2025 Argonne National Laboratory survey of 1,689 local emergency management agencies, more than half reported having one or no permanent full-time employees.[1] State directors surveyed by Deloitte and the National Emergency Management Association named budget constraints and a competitive job market (81% each) and a shortage of qualified candidates (64%) as their top workforce challenges.[2] At the same time, the planning workload keeps growing: emergency operations plans, hazard mitigation plans, THIRA/SPR submissions, grant justifications, and after-action reports, each carrying its own compliance requirements and cross-references.

Into this gap arrives generative artificial intelligence, and with it a common anxiety inside the profession: that AI will automate away the role of the planner.

This paper challenges that assumption. Emergency management is fundamentally a human endeavor rooted in relationship management, community trust, and multi-agency coordination. No model can sit across the table from a hospital administrator and negotiate a resource-sharing agreement. What AI can do, when properly governed, is absorb the administrative weight that currently keeps planners at their keyboards instead of in their communities.

The pages that follow outline a framework in which AI acts as a force multiplier inside the CPG 101 planning process, the risks that any public safety agency must manage before adoption, a human-in-the-loop governance model, and a practical three-phase roadmap an agency of any size can begin this quarter. The bottom line: agencies that treat AI as a drafting and analysis assistant, and reinvest the recovered hours in stakeholder engagement, will scale capability without surrendering judgment.

1. The Capacity Problem Is Already Here

The emergency management workforce conversation usually centers on response staffing. The quieter crisis is in preparedness. Local programs run by a single coordinator, or by an official wearing emergency management as a second or third hat, are the norm rather than the exception across much of the country.[1] The same Deloitte-NEMA study found state agencies struggling to recruit and retain the skills the mission now demands, from data analysis to grants management.[2] Even the federal layer is stretched: FEMA entered the 2025 hurricane season with roughly 12% of its incident management workforce available to deploy.[3]

The workload has not shrunk to match. A compliant emergency operations plan still requires hazard analysis, capability assessment, stakeholder coordination, annex development, and crosswalks against federal and state guidance. For a one-person shop, those hundreds of hours come directly out of the activities that determine whether a plan works when it matters: exercises, partnerships, and community engagement.

> *When a planner spends a hundred hours formatting annexes and cross-referencing compliance criteria, those are a hundred hours not spent building the coalition that will actually execute the plan.*

2. Data Is Not Capability: The Human Imperative

FEMA's Comprehensive Preparedness Guide (CPG) 101, now in Version 3.1, organizes emergency planning around a six-step process: form a collaborative planning team; understand the situation; determine goals and objectives; develop the plan; prepare, review, and approve the plan; and implement and maintain the plan.[4] Read the steps closely and a pattern emerges. Only two of the six are primarily analytical. The rest are coordination: assembling the right people, balancing stakeholder priorities, securing approval from elected officials, and keeping

the plan alive inside a community's planning cycles. CPG 101's whole-community doctrine makes the point explicit, directing planners to engage the public, private, and nonprofit sectors as partners rather than audiences.

Artificial intelligence excels at processing data, recognizing patterns, and synthesizing massive inputs. But a plan is only as good as the coalition built around it, and without people, it is just data. AI cannot negotiate a mutual aid agreement, read the political temperature of a county commission, or rebuild a neighborhood's trust after a flood. Human coordination adds the context that turns a synthesized hazard picture into an executable operational strategy.

The division of labor is therefore not a slogan but a design principle: **machines synthesize, humans contextualize, and the plan that results is both faster to produce and more grounded in the community it protects.**

3. A Lighter Lift Across the Six Steps

This direction is no longer speculative. FEMA's own published AI use case inventory includes the Planning Assistant for Resilient Communities (PARC), a generative AI tool intended to help SLTT governments draft hazard mitigation plan elements from well-researched public sources, and an AI-assisted resource portal to match SLTT decision-makers with recovery and resilience programs.[5] Plan-support AI is coming to the SLTT environment whether individual agencies prepare for it or not. The question is how to integrate it deliberately. Mapped against CPG 101, the opportunities concentrate where the work is most administrative:

CPG 101 Step	Where AI Assists	Where Humans Lead
Step 1: Form a Collaborative Planning Team	Maps stakeholder landscapes from public records; drafts engagement materials and meeting summaries.	Recruits the team, manages relationships, and resolves competing agency priorities.
Step 2: Understand the Situation	Ingests decades of hazard history, GIS layers, census shifts, and infrastructure data into a current risk snapshot.	Validates findings against local knowledge and lived community experience the data misses.
Step 3: Determine Goals & Objectives	Drafts candidate objectives traceable to the risk analysis; flags conflicts with existing plans.	Sets priorities, makes trade-offs, and aligns objectives with political and fiscal reality.
Step 4: Plan Development	Produces first-draft annexes; cross-references content against federal and state planning criteria in minutes.	Facilitates courses-of-action sessions; decides what the agency will actually commit to.
Step 5: Plan Preparation, Review & Approval	Runs compliance and consistency checks; assembles review packages and change logs.	Briefs leadership, negotiates revisions, and secures formal promulgation.
Step 6: Plan Implementation & Maintenance	Monitors data pipelines for changes that invalidate planning assumptions; drafts AAR content from exercise records.	Runs the exercises, owns corrective actions, and keeps partners engaged between updates.

The pattern holds across every step: AI compresses the time between raw information and a working draft, and the planner spends the recovered hours on the parts of the process no machine can perform.

4. Planning for Communities That Won't Sit Still

A static, five-year planning rewrite cycle was a reasonable convention when communities changed slowly. It is no longer sufficient for a county adding subdivisions, distribution centers, and data centers faster than its plans can be amended. Every new corridor of growth changes evacuation behavior, resource demand, and the population that needs to be warned.

AI-driven monitoring lets a small agency hold a dynamic posture that previously required a planning staff it could never afford. As new infrastructure comes online or demographics shift, continuously monitored data pipelines flag

how the change alters the local risk profile, and tools like FEMA's Resilience Analysis and Planning Tool (RAPT), referenced in CPG 101 Version 3.1 itself, already point in this direction.[4] The collaborative planning team can then adjust evacuation routes, resource positioning, and public information strategies proactively, on the community's timeline rather than the rewrite cycle's.

5. The Risks Nobody Should Skip

Thought leadership on AI that ignores its failure modes is marketing. Four risk categories deserve attention before any SLTT agency puts AI inside its planning process:

- **False confidence and confabulation.** Generative models produce fluent, plausible text that can be wrong. A fabricated shelter capacity or an invented mutual aid resource embedded in an EOP is not an inconvenience; it is an operational hazard. NIST's generative AI guidance identifies confabulation as a defining risk of the technology, which is precisely why outputs must be verified by someone who knows the jurisdiction.[6]
- **Data security and privacy.** Planning data includes critical infrastructure details, vulnerable population information, and personally identifiable information. Agencies must know where a vendor's model processes data, what it retains, and whether the arrangement satisfies applicable state and federal data handling requirements before any sensitive material enters a prompt.
- **Equity and bias.** Models trained on historical data inherit historical blind spots. Communities that were undercounted, underinsured, or underserved in past data will be underweighted in algorithmic risk pictures. The whole-community doctrine at the center of CPG 101 demands that human planners actively correct for this, not assume the model sees everyone.
- **Transparency and public accountability.** Emergency plans are public documents, and the analysis behind them may be subject to records requests and legal scrutiny. An agency must be able to explain how a conclusion was reached. 'The model said so' is not an answer that survives an after-action review or a courtroom.

None of these risks argues against adoption. They argue for governance. The NIST AI Risk Management Framework gives agencies a voluntary, widely adopted structure for doing exactly that, and federal policy now requires agencies to inventory AI use cases and apply minimum risk management practices to high-impact systems.[6] [7] SLTT agencies do not need to invent a governance model; they need to borrow one and enforce it.

6. Human-in-the-Loop by Design: The Sentinel Governance Model

Sentinel Resilience Partners applies a strict division of responsibility designed to protect public safety operations. The principle is simple: AI drafts and analyzes; the emergency manager decides and owns.

Capability	AI Responsibility	Human Responsibility
Hazard & Risk Ingestion	Rapid aggregation of multi-source GIS, historical incident, and demographic data.	Validation against localized history, community impacts, and ground truth.
Plan Drafting & Auditing	First-draft annex content; immediate compliance mapping against federal and state doctrine.	Strategic approval of operational goals, agency assignments, and final language.
Resource & Scenario Modeling	Predictive analysis of supply chains, asset allocation, and population movement.	Relationship-driven execution, mutual aid coordination, and real-world judgment.
Continuous Plan Maintenance	Monitoring data pipelines; flagging assumptions invalidated by community change.	Deciding what changes, briefing leadership, and re-promulgating the plan.

Every AI output passes through a verification layer before it touches an operational document, and nothing is published, promulgated, or briefed without a named human owner. The emergency manager remains the definitive strategic leader; the AI remains an assistant with a fast keyboard.

7. A Practical Adoption Roadmap

Most SLTT agencies cannot hire an AI officer, and they do not need to. A staged approach keeps risk proportional to experience:

Phase 1 — Govern First (Months 0–3)

Before any tool is procured, adopt a short written AI use policy: what data may and may not enter external systems, which outputs require verification, and how AI assistance is disclosed in published documents. The NIST AI RMF's four functions (Govern, Map, Measure, Manage) scale down surprisingly well to a two-page policy for a small agency.[6]

Phase 2 — Pilot the Low-Risk Lifts (Months 3–9)

Start where errors are cheap and visible: meeting summaries, draft after-action report sections, grant application boilerplate, plan formatting, and compliance crosswalks of existing plans. These tasks consume disproportionate staff hours, and a human reviewer catches mistakes before they propagate.

Phase 3 — Integrate into the Planning Cycle (Months 9–18)

With governance habits established, bring AI into the six-step process itself: hazard data synthesis in Step 2, draft annex development in Step 4, and continuous plan monitoring in Step 6. Measure success in one currency above all others: staff hours returned to stakeholder engagement, exercises, and community presence.

Conclusion

The question facing the profession is not whether AI will change emergency planning. Federal agencies are already building plan-support AI for the SLTT audience, and the technology is already in the office through every browser. The question is whether emergency managers will govern these tools deliberately or absorb them by accident.

The planners who thrive in the next decade will not be the ones who type the fastest. They never were. The profession's value has always lived in trust, coordination, and judgment, and those are precisely the assets AI cannot replicate and can only protect, by giving the people who hold them their time back.

> **About Sentinel Resilience Partners.** Sentinel Resilience Partners is a strategic advisory firm specializing in emergency management, crisis preparedness, continuity, and organizational resilience for public and private sector clients. Sentinel's ALIGN methodology applies the same doctrine-grounded discipline described in this paper, pairing CPG 101 and National Preparedness System rigor with modern analytical tooling under permanent human-in-the-loop governance. To discuss what this approach could look like in your jurisdiction, visit sentinelresiliencepartners.com.

References

1. Argonne National Laboratory, survey of 1,689 local emergency management agencies (July 2025), as reported in ProPublica's reporting on local emergency management staffing.
2. Deloitte Center for Government Insights & National Emergency Management Association, "National Risk Study: Changing Landscapes in State Emergency Management" (2025).
3. U.S. Government Accountability Office, "FEMA Disaster Workforce" (GAO-23-105663) and subsequent 2025 reporting on FEMA incident workforce availability.
4. Federal Emergency Management Agency, "Developing and Maintaining Emergency Operations Plans: Comprehensive Preparedness Guide (CPG) 101," Version 3.1.
5. U.S. Department of Homeland Security, Artificial Intelligence Use Case Inventory: FEMA — Planning Assistant for Resilient Communities (PARC); Recovery and Resilience Resource Portal.
6. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)" and the Generative Artificial Intelligence Profile (NIST AI 600-1).
7. Office of Management and Budget, Memorandum M-25-21, "Accelerating Federal Use of AI through Innovation, Governance, and Public Trust" (2025).