COALITION FOR INNOVATION

AI
Blueprint for the Future

Coalition for Innovation, supported by LG NOVA

Jami Diaz, Director Ecosystem Community & Startup Experience
William Barkis, Head of Grand Challenges & Ecosystem Development
Sokwoo Rhee, Executive Vice President, LG Electronics, Head, LG NOVA

Coalition for Innovation Co-Chairs

Alex Fang, CleanTech Chair
Sarah Ennis, AI Chair
Alfred Poor, HealthTech Chair

Authors

Adrien Abecassis, Johnny Aguirre, John Barton, Ann M. Marcus, Olivier Bacs, Taylor Black, Micah Boster, Mathilde Cerioli, Carolyn Eagen, Sarah Ennis, Annie Hanlon, Christina Lee Storm, Andrew Yongwoo Lim, Jess Loren, Refael Shamir, Svetlana Stotskaya

The views and opinions expressed in the chapters and case studies that follow are those of the authors and do not necessarily reflect the views or positions of any entities they represent.

Senior Editor, Alfred Poor
Editor, Jade Newton

October 2025

Preamble

The Coalition for Innovation is an initiative hosted by LG NOVA that creates the opportunity for innovators, entrepreneurs, and business leaders across sectors to come together to collaborate on important topics in technology to drive impact. The end goal: together we can leverage our collective knowledge to advance important work that drives positive impact in our communities and the world. The simple vision is that we can be stronger together and increase our individual and collective impact on the world through collaboration.

This “Blueprint for the Future” document (henceforth: “Blueprint”) defines a vision for the future through which technology innovation can improve the lives of people, their communities, and the planet. The goal is to lay out a vision and potentially provide the framework to start taking action in the areas of interest for the members of the Coalition. The chapters in this Blueprint are intended to be a “Big Tent” in which many diverse perspectives and interests and different approaches to impact can come together. Hence, the structure of the Blueprint is intended to be as inclusive as possible in which different chapters of the Blueprint focus on different topic areas, written by different authors with individual perspectives that may be less widely supported by the group.

Participation in the Coalition at large and authorship of the overall Blueprint document does not imply endorsement of the ideas of any specific chapter but rather acknowledges a contribution to the discussion and general engagement in the Coalition process that led to the publication of this Blueprint.

All contributors will be listed as “Authors” of the Blueprint in alphabetical order. The Co-Chairs for each Coalition will be listed as “Editors” also in alphabetical order. Authorship will include each individual author’s name along with optional title and optional organization at the author’s discretion.

Each chapter will list only the subset of participants that meaningfully contributed to that chapter. Authorship for chapters will be in rank order based on contribution: the first author(s) will have contributed the most, second author(s) second most, and so on. Equal contributions at each level will be listed as “Co-Authors”; if two or more authors contributed the most and contributed equally, they will be noted with an asterisk as “Co-First Authors”. If two authors contributed second-most and equally, they will be listed as “Co-Second Authors” and so on.

The Blueprint document itself, as the work of the group, is licensed under the Creative Commons Attribution 4.0 (aka “BY”) International License: https://creativecommons.org/licenses/by/4.0/. Because of our commitment to openness, you are free to share and adapt the Blueprint with attribution (as more fully described in the CC BY 4.0 license).

The Coalition is intended to be a community-driven activity and where possible governance will be by majority vote of each domain group. Specifically, each Coalition will decide which topics are included as chapters by majority vote of the group. The approach is intended to be inclusive so we will ask that topics be included unless they are considered by the majority to be significantly out of scope.

We intend for the document to reach a broad, international audience, including:

- People involved in the three technology domains: CleanTech, AI, and HealthTech
- Researchers from academic and private institutions
- Investors
- Students
- Policy creators at the corporate level and all levels of government

Chapter 12:
Agentic AI

Authors: Sarah Ennis, Taylor Black, Micah Boster, Ann M. Marcus

Introduction

Consider a scenario that plays out thousands of times daily in customer service centers: A customer calls about a billing discrepancy. They received a charge for a service they believe they canceled, but they're not sure when. The issue touches multiple systems including billing records, service activation logs, customer communications, and cancellation requests. A human agent must navigate between different databases, piece together the timeline, identify the root cause, apply appropriate credits or adjustments, update the customer's record, and send follow-up documentation.

Traditional AI can excel at individual components of this workflow such as analyzing billing data, generating explanations, or drafting customer communications. But it cannot autonomously orchestrate the entire resolution process. Each step requires a new prompt, a new context, and human oversight to connect the pieces. The customer waits while the agent manually bridges the gaps between AI-assisted tasks.

This gap between task-level AI assistance and end-to-end problem resolution represents one of the most significant limitations of current AI deployments. Organizations have invested heavily in AI tools that can summarize documents, generate content, or answer questions, yet find themselves still constrained by fundamentally human-driven workflows for complex, multi-step challenges.

Agentic AI represents a fundamental shift toward autonomous digital workers capable of independently managing these complex workflows from initiation to completion. Agentic systems interpret high-level goals, plan multi-step strategies, coordinate across tools and systems, and adapt in real time while maintaining appropriate oversight and control.

This transformation has profound business implications. Leading research organizations identify agentic AI as a top technology trend for 2025, while the market has grown from virtually nothing to $5.2 billion in 2024, with projections reaching $47 billion by 2030. Early adopters report $3-10 returns for every dollar invested, but more importantly, they're achieving operational capabilities that were previously impossible to automate.

Defining Agentic Behavior

Agentic AI systems exhibit four core characteristics that distinguish them from conventional AI applications, each representing a significant leap in autonomous capability.

Goal-oriented autonomy enables these systems to interpret high-level business objectives and independently determine the sequence of actions needed to achieve them. Consider the difference between asking traditional AI "What are our top customer complaints this month?" versus asking an agentic system "Improve customer satisfaction ratings." The traditional system provides data, while the agentic system analyzes complaint patterns, identifies root causes, researches solutions, proposes improvements, drafts implementation plans, and can even begin executing approved changes based on a single high-level directive.

Multi-step reasoning allows agentic systems to maintain context and adapt strategies across extended workflows that unfold over hours, days, or weeks. When a financial services company's agentic system detects unusual account activity, it doesn't just flag the transaction. It analyzes the customer's historical patterns, cross-references fraud databases, evaluates risk, determines appropriate responses, initiates security measures, prepares notifications, and schedules follow-ups while continuously monitoring for new signals that might change its assessment.

Dynamic tool integration represents perhaps the most transformative capability. Rather than being limited to pre-configured functions, agentic systems can discover, evaluate, and orchestrate whatever tools they need based on situational requirements. A research agent investigating market trends might seamlessly transition from web searches to database queries to statistical analysis tools to document generation platforms, selecting and combining tools in real-time based on the evolving information needs of its investigation.

Adaptive learning enables agentic systems to modify their behavior based on results and feedback, creating continuous improvement cycles. Unlike traditional AI that follows predetermined patterns, agentic systems evaluate their performance, identify failure points, and adjust their approaches. A content generation agent that notices that certain article types receive higher engagement will gradually shift its strategy, testing new approaches and incorporating successful patterns into its standard operating procedures.

The cumulative effect of these capabilities transforms AI from a sophisticated assistant that requires constant direction into an autonomous worker capable of managing complex business processes independently. This shift enables organizations to automate not just individual tasks, but entire workflows that previously required human judgment and coordination.

We explore real-world examples across public, private, and community domains in Section 5.

Core Architectures

Four foundational technologies have converged to make agentic AI practical for enterprise deployment. Understanding these building blocks helps explain both current capabilities and the technical challenges that remain unsolved. Each represents significant engineering advances, but also areas where "this sounds like automation we already have" skepticism is common and often misplaced.

Retrieval-Augmented Generation

Large language models face two fundamental limitations that constrain their business value: training cutoffs that create knowledge gaps, and complete ignorance of organization-specific information. A model trained on public internet data through 2023 knows nothing about your company's products, processes, customers, or internal knowledge base. More critically, it cannot access real-time information about inventory levels, customer interactions, regulatory changes, or market conditions that drive business decisions.

Retrieval-Augmented Generation (RAG) transforms AI from generic assistants into specialized business intelligence systems by enabling dynamic access to proprietary and current information. Rather than relying solely on training data, RAG systems actively search your databases, documents, customer records, and external sources to find relevant context before generating responses.

The technical implementation involves several sophisticated challenges. Proprietary data exists in diverse formats such as structured databases, unstructured documents, real-time feeds, legacy systems with inconsistent schemas. RAG systems must parse these varied sources, understand semantic relationships across different data types, maintain data lineage for compliance, and ensure security boundaries are respected. Real-time updates add another layer of complexity, as systems must balance freshness with computational efficiency while handling concurrent access to live data sources.

Consider a pharmaceutical company deploying RAG for regulatory compliance. The system must access clinical trial databases, FDA correspondence, internal protocol documents, published research, and regulatory filing histories, while also understanding the temporal relationships and approval dependencies that determine what information is relevant for specific queries. The technical challenge lies not just in searching these sources, but in understanding how different types of evidence combine to support regulatory decisions.

Advanced RAG implementations achieve 90% accuracy in data extraction across various formats while processing millions of documents monthly. These systems don't just retrieve information; they evaluate source authority, identify potential conflicts between sources, synthesize findings across multiple documents, and provide transparent attribution for verification. This enables true domain specialization where AI agents become organizational knowledge experts, converting broad intelligence into precise, context-aware decision-making capabilities.

Model Context Protocol and Tool Integration

One of the biggest barriers to deploying AI agents has been integration complexity. Without standardized protocols, connecting M AI agents to N external tools requires building M×N custom integrations, where each agent needs a separate connection to each tool, database, or system. This creates an exponential scaling problem as organizations add more agents and tools.

The Model Context Protocol (MCP) addresses this by creating a universal communication layer between AI agents and external resources. Based on the proven Language Server Protocol from software development, MCP transforms the complex M×N integration problem into a manageable M+N architecture. Instead of each agent requiring custom connections to every tool, agents connect to MCP servers that provide standardized access to external resources.

This means organizations can deploy standardized MCP servers that any compatible AI agent can utilize, eliminating the need for custom integrations. Industry adoption has been rapid, with Anthropic integrating MCP natively into Claude Desktop, OpenAI announcing MCP support for ChatGPT and their Agents SDK, and major platforms including Google's Gemini and Microsoft's frameworks following suit.

MCP standardization is crucial for enterprise deployment because it enables universal connectivity to any external system through a single protocol, dynamic tool discovery allowing agents to find and use new capabilities without code changes, and standardized security models with consistent permission and consent frameworks across all integrations.

Multi-Agent System Architectures

Rather than building monolithic AI systems that attempt to handle all tasks, multi-agent architectures deploy teams of specialized agents, each optimized and fine-tuned for specific capabilities while coordinating through sophisticated communication protocols.

The key insight is agent specialization, which involves creating AI agents specifically optimized for particular roles through specialized prompting, training data, or configuration. Just as human software teams benefit from having dedicated system architects, developers, QA specialists, and UX designers, AI agent teams can deploy specialists optimized for different aspects of complex workflows.

Enterprise orchestration frameworks exemplify this approach with asynchronous, event-driven architectures that enable natural language coordination between specialized agents. A software development workflow might deploy a system architecture agent optimized for technical planning, a code generation agent fine-tuned for specific programming languages, a quality assurance agent specialized in testing methodologies, and a UI/UX agent focused on user experience principles.

Multi-agent orchestration and management systems represent one of the most active areas of current development in agentic AI. While early implementations show promising results, the coordination mechanisms, communication protocols, and error handling systems are rapidly evolving. Organizations should expect dramatic improvements in maturity, reliability, and ease of deployment over the next 12-18 months as these frameworks advance.

Real-world implementations demonstrate both the potential and current limitations. Organizations have deployed multi-agent systems that reduced software requirements writing from weeks to days by orchestrating specialized agents for user story creation, technical analysis, and test plan documentation. However, performance benchmarks reveal that while specialized agent teams achieve impressive results on domain-specific tasks, general-purpose coordination still faces challenges with complex, multi-step workflows.

The core challenge becomes clear when you consider something as simple as pizza delivery. When the delivery person arrives at your door, a complex interaction unfolds, greeting, confirming the order, processing payment, and parting ways. These interactions flow naturally because humans have evolved sophisticated social protocols over millennia. AI agents, by contrast, are brilliant specialists trapped in digital isolation. They can analyze data, generate code, or write content with remarkable skill, but they have no innate understanding of how to coordinate with each other. A coding agent doesn't know when to hand off work to a testing agent, or how to communicate that it has encountered an error, or what to do when another agent goes offline mid-task. Much of current multi-agent development focuses on solving this fundamental interaction problem by teaching AI agents the basic social skills that allow them to work together rather than simply work in parallel.

But coordination is only part of the challenge. As AI systems grow more autonomous, the bottleneck is no longer capability but oversight. Traditional human-in-the-loop models don’t scale. The next leap forward is agentic AI that governs itself through internal red-teaming. Instead of relying solely on external human evaluators, specialized agents act as internal auditors, rigorously testing outputs for logic errors, hallucinations, and compliance gaps before results move downstream. This creates a dynamic ecosystem of peer review, where agents challenge, refine, and validate each other’s work. Such self-auditing architectures establish checks and balances that enable safe autonomy at scale, reducing reliance on human gatekeepers while increasing robustness, adaptability, and trustworthiness. Of course, evaluator agents are not immune to flaws; they too can misjudge, hallucinate, or become misaligned. That’s why recursive oversight is essential, with higher-order agents or consensus mechanisms monitoring the monitors and creating a layered defense against failure. The goal is not perfection but resilience through distributed accountability.

Planning and Reasoning Frameworks

Traditional AI systems respond to immediate prompts but struggle with complex, multi-step challenges that require strategic thinking. Advanced planning frameworks transform reactive systems into strategic thinkers capable of sophisticated workflow orchestration.

A crucial capability that distinguishes agentic systems is their ability to combine AI-driven reasoning with deterministic operations within the same workflow. An agent might use AI to analyze customer feedback data and identify patterns, then execute precise SQL queries to retrieve specific customer records, perform mathematical calculations on the results, and finally use AI again to generate personalized recommendations. This hybrid approach leverages AI's interpretive capabilities alongside the reliability and precision of traditional computational methods.

Planning frameworks implement sophisticated decision-making processes where agents analyze situations, consider multiple approaches, execute both AI-driven and deterministic actions, process results, and adjust strategies dynamically. The key innovation lies in intelligent workflow orchestration, which means knowing when to use AI for interpretation and creativity versus when to use deterministic processes for precision and reliability.

Performance results demonstrate these frameworks' effectiveness: 92.7% accuracy on programming tasks and 75.9% average scores on complex navigation challenges. The innovation lies in self-reflection mechanisms that enable agents to evaluate their own decision quality and learn from mistakes, creating a continuous improvement cycle essential for autonomous systems.

Hierarchical planning approaches enable agents to operate at multiple abstraction levels simultaneously by applying AI for high-level strategic thinking while executing precise deterministic operations for specific tasks. An agent might use AI reasoning to determine that a customer complaint requires account adjustment, then execute deterministic database updates to implement the change, and finally use AI again to craft an appropriate customer communication.

This combination of AI flexibility with deterministic reliability makes agentic systems far more powerful and trustworthy for enterprise applications, where both creative problem-solving and precise execution are essential for business-critical workflows.

Implementation Approaches

The choice between open-source frameworks and commercial platforms for agentic AI is a fundamental strategic decision, shaping an organization's long-term flexibility, costs, and adaptability to evolving AI capabilities. Skeptics often dismiss this as irrelevant, arguing "AI is AI," but this overlooks the profound technical complexities and strategic implications of deployment. Many organizations fail to grasp what happens when these systems inevitably encounter edge cases, break, or require modification for changing business needs, highlighting the critical importance of selecting the right implementation approach from the outset.

Open-Source Frameworks

LangChain dominates the open-source landscape with over 100,000 GitHub stars and more than one million monthly downloads. The framework's comprehensive ecosystem includes LangGraph for multi-agent orchestration, LangSmith for observability, and extensive integrations across the AI development stack. Major implementations serve tens of millions of users with significantly faster resolution times, demonstrating production-ready capabilities.

CrewAI has emerged as a preferred choice for teams new to agentic AI, emphasizing simplicity and role-playing agent interactions. The framework enables complex multi-agent workflows with minimal code, making it ideal for rapid prototyping and straightforward collaborative systems.

Microsoft's AutoGen targets enterprise environments with battle-tested reliability and sophisticated conversation-based coordination. The framework's asynchronous architecture and advanced error handling make it suitable for production environments where reliability is paramount.

Open-source advantages include complete customization control, transparency in operations, cost efficiency for organizations with technical expertise, and freedom from vendor lock-in. Organizations can modify frameworks to meet specific requirements, understand exactly how their AI systems operate, and avoid dependencies on external providers.

Implementation challenges include steep learning curves, frequent updates requiring ongoing maintenance, limited enterprise support, and the need for significant internal technical expertise. Organizations must invest in dedicated teams to manage, customize, and maintain these frameworks effectively.

Commercial Orchestration Platforms

Commercial agentic AI platforms provide comprehensive orchestration environments that handle the complexity of multi-agent coordination, tool integration, and workflow management through managed services. These platforms focus on business process automation rather than individual model capabilities.

n8n represents a leading workflow automation platform that has evolved to support AI agent orchestration. The platform provides visual workflow builders, extensive integrations with business tools, and sophisticated error handling for complex multi-step processes. Its strength lies in enabling non-technical users to create sophisticated agent workflows while maintaining enterprise-grade reliability and monitoring.

The no-code/low-code automation space has rapidly expanded to include agentic AI capabilities. Platforms like Make (formerly Integromat) and newer entrants like Gumloop provide visual workflow designers specifically optimized for AI agent coordination. These platforms democratize agentic AI by allowing business users to create complex multi-agent workflows without programming expertise, often at significantly lower costs than enterprise solutions. However, they may lack the advanced error handling and enterprise governance features required for mission-critical applications.

Full-service development platforms represent another emerging category. Replit's AI-powered development environment enables rapid prototyping and deployment of agentic applications, while platforms like Loveable focus on end-to-end AI application development with built-in agent orchestration capabilities. These platforms blur the line between development tools and deployment environments, offering integrated solutions for organizations that want to build custom agentic applications without extensive infrastructure investment.

Microsoft's Power Platform, including Power Automate and Copilot Studio, offers deep integration with Microsoft's business ecosystem. The platform excels at connecting AI agents with existing productivity workflows, providing seamless handoffs between human and AI workers within familiar business applications. This makes it particularly valuable for organizations already invested in Microsoft's technology stack.

Specialist application platforms represent another category, exemplified by Salesforce's Agent Force, which focuses specifically on customer relationship management and sales process automation. Rather than providing general-purpose orchestration, these platforms offer deep domain expertise within their specific business functions. Agent Force understands customer journey orchestration, maintains context across multiple touchpoints, and integrates natively with CRM data and sales processes. This approach offers significant advantages for organizations whose agentic AI needs align with the platform's specialization, but limits flexibility for use cases outside that domain.

Enterprise workflow platforms like UiPath and Automation Anywhere have expanded beyond traditional RPA to include AI agent capabilities, offering the advantage of integrating agentic AI with existing automation infrastructure. These platforms excel in environments where AI agents need to work alongside traditional automated processes.

These commercial solutions offer distinct advantages: immediate deployment capabilities, visual workflow designers accessible to business users, managed infrastructure with automatic scaling, enterprise security and compliance features, and professional support with service level agreements. However, organizations must weigh these benefits against vendor lock-in risks, higher long-term costs, and reduced customization flexibility compared to open-source alternatives.

Selection criteria should prioritize alignment with organizational capabilities and integration requirements: no-code platforms like Make and Gumloop for rapid deployment by business users, full-service platforms like Replit and Loveable for custom application development, Microsoft for productivity-focused environments, specialist platforms like Salesforce for domain-specific applications, and enterprise RPA platforms for organizations with existing automation infrastructure.

Commercial advantages include immediate deployment capabilities, professional support and service level agreements, managed infrastructure with automatic updates, enterprise security and compliance features, and reliability guarantees suitable for customer-facing applications.

Considerations include higher long-term costs, potential vendor lock-in, limited customization options, and dependency on external providers for critical business functions.

Hybrid Strategies

Leading organizations increasingly adopt hybrid strategies that leverage both open source and commercial solutions strategically. Development teams use open-source frameworks for research, prototyping, and internal applications while deploying commercial services for customer-facing systems requiring reliability and support.

Multi-vendor approaches combine different commercial services based on specific strengths, with some providers focused on complex reasoning, others on safety-critical applications, and others on multimodal tasks. This strategy mitigates vendor risk while optimizing capabilities for different use cases.

Cost analysis reveals that hybrid approaches typically achieve 20-40% cost savings compared to pure commercial solutions while maintaining enterprise-grade capabilities. Open-source development costs range from $20,000 to $500,000+ depending on complexity, while commercial solutions cost $100-5,000 monthly for standard implementations.

Success factors for hybrid implementations include standardized infrastructure using universal protocols, unified governance frameworks, strong internal technical capabilities, and clear decision criteria for when to use each approach.

Real-World Applications Across Sectors and Risks/Benefits

Agentic AI is already being piloted across public, private, and nonprofit sectors. From emergency evacuations to internal compliance agents, the spectrum of uses is rapidly expanding. The following table highlights where AI agents are beginning to take root and the key functions they perform.

AI Agentic Use Examples Across Different Sectors

Sector	Example Agent Function
Public	Evacuation logistics, permit navigation, civic updates
Private	Compliance audits, internal project agents, client service
Community	Outreach, translation, mutual aid coordination
Commercial	Travel planners, smart shopping, home automation

Public Sector Uses (Government / Infrastructure)

Use Case	Description
Emergency Evacuation Coordination	AI agents manage logistics for evacuating vulnerable populations during disasters, as demonstrated in the senior evacuation model in Appendix A
Digital Permit & Licensing Agents	Agents guide residents through applications for permits (e.g., building, business, event), auto-filling and submitting forms.
Public Transportation	Agents help commuters navigate multi-modal transit systems in real time,

Routing Assistants	suggesting accessible or low-cost options.
Civic Engagement Bots	Agents summarize city council meetings, propose meeting agendas, or alert residents to decisions affecting their neighborhood.
Climate Risk Notification Agents	Personalized early-warning systems that notify individuals of local environmental risks based on location and profile.

Private Sector Uses (Corporate / Institutional)

Use Case	Description
Enterprise Workflow Optimization	AI agents monitor project progress, flag bottlenecks, and suggest next steps or staffing reallocation in real time.
Smart Scheduling Assistants	Agents coordinate across internal calendars and meeting goals to arrange cross-team availability or escalate urgent requests.
Compliance Monitoring Agents	Track evolving regulations and assess company compliance gaps, especially in data privacy, ESG, or workplace safety.
Customer Service Delegation	Empowered agents handle multi-channel customer interactions, escalate only complex cases to human staff.
AI for IT Support	Autonomous agents resolve software configuration issues, patch systems, or auto-escalate based on system anomalies.

Community Benefit / Nonprofit Uses

Use Case	Description
Disaster Response and Recovery Agents	Coordinate supply distribution (water, food, shelter), volunteer deployment, and damage assessment.
Elder Outreach & Wellness Check-ins	Voice-capable agents call isolated seniors regularly, assess their mood or needs, and escalate alerts as necessary.

Neighborhood Improvement Agents	Automate surveys to gather community feedback, propose mini-projects (e.g., tree planting, sidewalk repairs), and track progress.
Language Access for Immigrants	Translation agents assist non-English speakers in accessing healthcare, housing, or legal services.
Civic Literacy Bots	Agents explain ballot measures, voter registration steps, or public program eligibility in plain language.

Commercial / Consumer-Facing Uses

Use Case	Description
Personal Shopping Agents	AI agents curate products based on user needs, search across platforms, compare pricing, and place orders.
Travel Booking & Rescheduling	Agents auto-plan travel (flights, hotels, transport) based on constraints like budget, loyalty points, and accessibility.
Home Energy Optimization	Agents learn usage patterns and adjust HVAC, lighting, and appliances to lower bills and carbon footprint.
Gig Worker Schedulers	Agents manage freelance jobs, match workers with demand, and optimize routes or shifts.
Education / Tutoring Agents	Personalized AI tutors support students in learning at their pace, flag gaps, and adjust teaching methods accordingly.

Potential Risks Associated with Using Agentic AI in Various Domains:

Domain	Key Risks
Public	Bias, accountability gaps, data misuse, cyber threats
Private	Oversight loss, security leaks, job displacement
Community	Consent, equity, miscommunication, loss of trust
Commercial	Privacy erosion, manipulation, financial errors

Mitigation Considerations

To reduce the risks of agentic AI deployment, organizations should implement:

Mitigation Strategy	Purpose
Human-in-the-loop oversight	Maintains accountability and decision control
Ethical review panels or audits	Evaluates fairness, safety, and unintended outcomes
Community co-design	Ensures inclusivity and local relevance
Privacy and consent safeguards	Protects sensitive data and user autonomy
Monitoring and feedback loops	Detects errors, drift, or unintended behaviors early
Audits and adjustment cadence	Enables structured iteration and performance tuning

Getting Started with Agentic AI

Organizations looking to adopt agentic AI should begin with low-risk, internal workflows such as compliance monitoring, project tracking, or IT automation. Start small:

- **Pilot in controlled environments** where outputs can be safely evaluated.

- **Map existing toolchains** to identify integration gaps or friction points.
- **Use evaluator agents** to red-team outputs before broader rollout.
- **Define fail-safes** for critical steps where accuracy or accountability is key.
- **Track performance** and iterate with clear metrics tied to cost, speed, or quality gains.

Starting this way builds confidence, reveals edge cases early, and creates a foundation for scaling agentic systems responsibly.

Appendix I: Agentic AI in Disaster Response

The following extended case study illustrates how Agentic AI can be deployed in a complex, high-stakes, public-sector context: disaster response for vulnerable populations.

Climate change is dramatically increasing the frequency, intensity, and unpredictability of disasters such as wildfires, floods, heat waves, earthquakes, and tsunamis. These pose heightened risks for elderly and disabled populations. These individuals are disproportionately affected by delayed or inaccessible evacuation efforts, yet most municipalities across the U.S. (and globally) remain woefully underprepared to respond effectively.

Agentic AI -- which refers to autonomous, goal-driven software systems – offers a promising solution to bridge the gap between emergency response plans and real-time operational coordination. These intelligent agents can be deployed to ensure timely, adaptive, and inclusive evacuation strategies by performing the following functions:

Key Agentic AI Functions in Evacuation Coordination

Proactive Outreach and Needs Assessment

- AI agents can identify and reach out to registered seniors, disabled individuals, or others on medical alert or community watchlists.
- Using phone, SMS, or voice interfaces, agents can assess evacuation status, transportation needs, medical dependencies, or mobility constraints.

Dynamic Transportation Coordination

- Agents can tap into multi-modal transportation networks such as public buses, commercial ride-shares (such as Uber WAV), paratransit services, non-emergency medical transport, accessible taxis, and vetted volunteer drivers.
- AI agents dynamically match evacuees with appropriate vehicle types, prioritizing mobility needs, proximity, and urgency.

Multi-Agency Communication and Dispatch

- AI agents can serve as intermediaries between emergency command centers, transportation providers, shelters, and health services, ensuring unified situational awareness.
- AI agents are capable of real-time updates, rerouting, and reassignment as hazards evolve (e.g., wildfire direction changes or road closures).

Support for Caregivers and Families

- AI agents can notify designated caregivers or family members of the individual's status and whereabouts during transit.
- AI agents can also act as virtual assistants for self-advocating seniors, enabling voice-based check-ins or confirmations.

The San Leandro Senior Evacuation Project by WeAccel

WeAccel is actively developing a proof-of-concept senior evacuation model in San Leandro, California, integrating Agentic AI to:

- Establish a senior registry and risk map that includes mobility status, medical equipment needs, language preferences, and household situation,
- Coordinate with municipal emergency planners, transportation operators, and senior service organizations, and
- Pilot an AI-driven outreach and routing system that can operate with limited broadband or SMS-only infrastructure, which is crucial for underserved or tech-limited seniors.

This project aims to prototype a replicable framework for other cities and contribute to a resilience network that centers the most vulnerable in disaster planning.

Why It Matters

Without action, emergency events exacerbated by climate change will continue to result in preventable deaths and suffering among the elderly and disabled, particularly those who live alone, lack Internet access, or have limited ability to speak or understand English.

Agentic AI systems can dramatically reduce coordination delays, optimize resource use, and ensure no one is left behind, especially when these systems are built with community input, equity considerations, and redundancy planning in mind.

Stakeholder Participation & Required Data

To create and coordinate a system of AI agents that supports emergency evacuation for seniors and disabled individuals, it would be necessary to identify a wide range of stakeholders and data sources.

Below is a breakdown of both, organized by functional role and data dependencies.

Key Stakeholders

Public Sector & Emergency Management

- City and County Emergency Services Departments: Responsible for evacuation plans, EOCs (Emergency Operations Centers), alert systems
- Fire, Police, EMS: Need real-time access to evacuation routes and special needs populations
- Public Health Departments: Provide insight into medical vulnerabilities, home care needs, oxygen/electricity dependence
- Transportation Agencies: Coordinate buses, paratransit, and detours during emergencies

Community-Based Organizations (CBOs)

- Senior Centers & Aging Services Providers: Maintain contact lists, care plans, and wellness check routines
- Disability Rights Organizations: Ensure accessibility and advocate for inclusion in planning and execution
- Faith-Based and Mutual Aid Groups: Provide local trust and human support for outreach, ride-alongs, and wellness checks

Public, Private & Commercial Transport Providers

- City Vehicles
- Public Transportation Vehicles (e.g. AC Transit for Alameda County)
- Paratransit Services
- Ride-hailing Companies (e.g., Uber WAV, Lyft Access)
- Medical Transport Providers
- Charter or Shuttle Companies
- Taxi Services
- Volunteers with Registered Vehicles

Technology & Infrastructure Partners

- Telecom Providers: Enable SMS/voice connectivity and geolocation services
- AI Developers / Agentic AI Platforms: Build, train, and deploy AI agents capable of autonomous coordination, outreach, routing, and translation

- Data Integration Vendors: Handle cross-agency data aggregation, privacy, and interoperability
- Mapping & Navigation Tools: Enable real-time routing, congestion detection, and road hazard data

Funders & Oversight Bodies

- Local, State, and Federal Grant Authorities (FEMA, HUD, state emergency or aging offices): Provide funding for technology pilots, infrastructure, and resilience programs, while requiring compliance with emergency management standards.
- AARP and Aging Advocacy Organizations: Offer funding and legitimacy for senior-focused solutions, ensuring alignment with national aging and disability priorities.
- Foundations (e.g., Knight Foundation, Robert Wood Johnson Foundation): Support innovation, community-based pilots, and equity-focused approaches.
- Academic Research Partners: Evaluate system performance, test for bias, and strengthen models with evidence-based methods and community input.

Critical Data Requirements

Individual-Level Data (with consent or emergency-use authorization)

Data Type	Source / Provider	Notes
Name, Age, Address, Contact Info	Senior registries, utility bills, 911 databases	May require aggregation
Mobility Status (e.g., wheelchair)	CBOs, Health Departments	Includes care dependencies
Medical Needs (e.g., oxygen, meds)	Public Health, Home Health Agencies	Privacy-protected
Language Preference	Registries, CBO intakes	Enables multi-language AI
Household Composition	CBO intakes, utility records	Flags additional residents needing support

Agentic AI can transform emergency response for seniors through predictive monitoring, rapid alerting, and personalized care coordination. While these systems excel at connecting seniors with help in medical or facility-based emergencies, development continues towards fully autonomous, AI-driven platforms that directly match seniors seeking evacuation with providers during mass emergencies. Current approaches may employ AI to support and inform human responders who can execute ride arrangements. However, more automated approaches are likely to become more prevalent as deployments scale and reliability improves.

Some other examples of similar solutions include:

Austin, TX, Vulnerable Population Registries

Austin operates a Medically Vulnerable Registry run through Austin Energy and the broader State of Texas Emergency Assistance Registry (STEAR). These registries collect data on medically fragile persons to inform emergency planning, including evacuation, but without real-time AI-enabled coordination or multi-modal transport automation. While these systems improve situational awareness, there’s no evidence yet of integrated AI agents dynamically matching registrants to transport in real time during an event.

Japan Post-Tsunami Robotic & AI Tools

> Japan has a long history deploying rescue robots – including the snake-like Quince and tracked T-52 Enryu – in earthquake and tsunami zones to aid shelter access or debris-clearing operations. More recently, systems such as *Spectee Pro* use AI to analyze social media, weather, and satellite images to enhance situational awareness during disasters, but still focus on information gathering and shelter logistics, not on automated transport coordination. The city of Rikuzentakata, for example, launched automated calls to registered residents to check evacuation status, but again this is contact-based outreach without full AI-driven multimodal transit integration.

Even with these early efforts, the opportunity remains to develop a solution that fully integrates agentic AI, multi-source transportation coordination, and senior-centric mobility and accessibility needs. See WeAccel.io for more information on this project.

Appendix II: Key Federal & Regional Programs Supporting AI in Emergency Response

Federal Programs

Department of Homeland Security (DHS) AI Pilots

Scope: Nationwide

Focus: Safe and secure AI deployment as part of the AI Executive Order

Status: Phase 1 complete; AI Corps hired

Contact: DHS Science & Technology Division

Learn more: https://www.dhs.gov/science-and-technology/

FEMA AI Use Cases

Scope: Nationwide

Program: FEMA contributes to DHS's AI Use Case Inventory

Focus: Emergency management AI exploration

Contact: FEMA HQ via DHS

Access: https://www.dhs.gov/science-and-technology/

Regional and State Programs

Miami-Dade Emergency & Evacuation Assistance Program (EEAP)

- **Location**: Miami-Dade County, FL

Function: Helps residents with medical/special needs evacuate safely

Services: Specialized transportation coordination

Contact: Miami-Dade County Services

Website: https://www.miamidade.gov/global/service.page

Florida State Emergency Resources

Scope: Statewide

Services: Emergency support and health guidance

Public Hotline: 800-342-3557

Website: https://www.floridahealth.gov/about/emergency.html

Private Sector Solution

Prepared911 AI Platform

Type: Commercial, end-to-end AI emergency response system

Function: AI assistance integrated across the emergency response lifecycle

Pilot Opportunities: Available for agencies and municipalities

Website: https://www.prepared911.com/

References:

The Inflection Point: Agentic AI in the Evolution of Security and Risk Management – *Crisis24*
https://www.crisis24.com/articles/the-inflection-point-agentic-ai-in-the-evolution-of-security-and-risk-management

Agentic AI in Disaster Management and Emergency Response – *DigitalDefynd*
https://digitaldefynd.com/IQ/agentic-ai-in-disaster-management-and-emergency-response

Multi-Agent Systems in Disaster Management – *SmythOS*
https://smythos.com/ai-agents/multi-agent-systems/multi-agent-systems-in-disaster-management

How AI Boosts Emergency Response Times – *EyewatchLive*
https://eyewatchlive.com/news/how-ai-boosts-emergency-response-times

Innovations in Personal Emergency Response Systems for the Elderly – *Ball State Daily News*
https://www.ballstatedaily.com/article/2025/02/innovations-in-personal-emergency-response-systems-for-the-elderly

Agentic AI in Home Care – *AutomationEdge*
https://automationedge.com/home-health-care-automation/blogs/agentic-ai-in-home-care

Microsoft GraphRAG GitHub Repository
https://github.com/microsoft/graphrag

Anthropic Documentation – MCP Overview
https://docs.anthropic.com/en/docs/agents-and-tools/mcp

The Pragmatic Engineer Newsletter – MCP Deep Dive
https://newsletter.pragmaticengineer.com/p/mcp

Japan's Use of AI and Robotics in Disaster Response
https://www.researchgate.net/publication/4181423_Rescue_Robots_and_Systems_in_Japanhttps://spectrum.ieee.org/japan-earthquake-robots-help-search-for-survivors

https://www.preventionweb.net/news/japan-firms-look-ai-bolster-disaster-prevention-and-mitigation

Austin Energy Medically Vulnerable Registry Audit Report (2024)
https://www.austintexas.gov/sites/default/files/files/Auditor/Audit_Reports/Austin_Energy_Medically_Vulnerable_Registry_March_2024.pdf

Author (In order of contribution)

Sarah Ennis, Co-Founder and Advisor of AgentsGEO.ai
Sarah Ennis is a Fortune 500 trusted advisor specializing in advanced technology innovation, with over two decades of experience leading groundbreaking AI solutions at scale. Globally recognized for her expertise in artificial intelligence, she designs and implements bespoke emerging technology products across industries. She is also the co-founder and advisor of AgentsGEO.ai, a patent-pending

platform that helps brands monitor and improve their visibility in the AI ecosystem and deploy AI agents, ensuring they are discoverable and recommended by tools like ChatGPT, Gemini, and others through its proprietary GEOScorer™ technology. In addition, Sarah contributes part-time to Northeastern University's Master of Digital Media programs in AI, preparing the next generation of technologists and creative leaders. Her work bridges Silicon Valley innovation with global impact, and she is a distinguished member of the American Society for AI and contributor to the OpenAI Forum.

Taylor Black, Director AI & Venture Ecosystems, Microsoft
Taylor Black is Director of AI & Venture Ecosystems in Microsoft's Office of the CTO, where he designs and leads cross-company initiatives that integrate innovation, product development, and community engagement. With 19+ years of experience launching and scaling ventures across enterprise, deep tech, and social ecosystems, he brings a multidisciplinary background as a developer, educator, lawyer, entrepreneur, and venture builder. He mentors and invests in early-stage startups through networks such as Conduit Venture Labs and Fizzy Ventures. Taylor also helps shape Catholic University of America's new institute at the intersection of AI, innovation, and human flourishing.

Micah Boster, Principal, Nighthawk Advisors
Micah Boster is the founder and Principal at Nighthawk Advisors, where he works with early-stage technology companies on execution, AI strategy, and positioning. Previously, he spent eight years at Google and over a decade as an executive at several NYC-based startups. He holds a BS in Symbolic Systems from Stanford and an MBA from INSEAD.

Ann M. Marcus, Director, Ethical Tech & Communications, WeAccel
Ann M. Marcus is a Sonoma-raised, Portland-based communications strategist and ethical technology analyst focused on smart cities, community resilience, and public-interest innovation. She leads the Marcus Consulting Group and serves as director of ethical technology and communications at WeAccel.io, a public-good venture advancing mobility, communications, and energy solutions for communities. Ann has advised public and private organizations—including Cisco, the City of San Leandro, Nikon, AT&T, and InfoWorld—on trust-based data exchange, digital public infrastructure, resilience strategy, AI and more. Her current projects include a California senior evacuation program, a Portland robotics hub, and digital energy resource initiatives with utilities in Portland and the Bay Area.

For more information about the Coalition for Innovation,
including how you can get involved, please visit coalitionforinnovation.com.

View the Next Chapter