ATARC

Agentic AI Working Group

FEBRUARY 2026

A Roadmap for Success

Building a Comprehensive AI Adoption Playbook

Acknowledgements

Frank Indiviglio, NOAA
Howard Rosen, Nova Insights Corp, HIMSS
Ferdous Khan, FEMA

A structured and comprehensive AI Adoption Playbook is critical for any organization aiming to harness the transformative power of artificial intelligence. This playbook should serve as a practical guide, navigating the complexities of AI integration from executive buy-in to equitable implementation and scalable growth. The following expanded outline, enriched with real-world examples and best practices, provides a robust framework for building such a playbook.

McKinsey reports that only 30% of digital transformation projects are successful, meaning that 70% of these projects failed (and these were using technologies that were widely understood). The Clinger-Cohen Act directs Government organizations to use an Architectural process for IT acquisitions. This policy was created to prevent the Government from repeating technological disasters it experienced decades ago. The first step is to understand what the business processes are. A good place to start is with the organization's Strategic Plan. Existing systems automate the processes identified in the Strategic Plan. Mapping how existing systems support the mission is referred to as the AS-IS Architecture. The AS-IS Architecture is examined to determine if there are any places where the use of AI can improve these processes. This requires two things:

1. A knowledge of the kinds of things AI can be used for.

2. A data inventory. (A frequent cause of failures in AI projects is that the data is either not available or is inappropriate for the kind of AI.) Based on this, a TO-BE architecture is developed. In some cases, a new functionality, such as an enterprise search capability implemented by a RAG (Retrieval-Augmented Generation) is envisioned. In these cases, a Proof of Concept (PoC) or an MVP "Minimum Viable Product" approach can be useful because it can provide valuable services first, with enhancements and improvements following. In any case, once a TO-BE architecture is created, an Implementation Plan, with timelines and estimates of costs and potential returns on investment can be developed.

Foundation Level (Months 1-6): Building the Core

This initial phase focuses on establishing the essential building blocks for safe, responsible, and effective AI adoption. It's about getting the basic infrastructure, policies, and understanding in place.

1. Cultivating a Generative AI-Ready Culture

To successfully integrate and leverage Generative AI (GenAI) across an organization, the first step is to **actively cultivate a culture of experimentation and psychological safety**. Leaders must clearly communicate that GenAI is a tool for augmentation, not replacement, framing its introduction as a positive opportunity for upskilling and increased efficiency, rather than a threat. This involves dedicating resources for exploratory projects, creating sandbox environments where employees can safely test different tools and prompts without fear of failure or negative repercussions, and rewarding early adopters who share successful use cases. Furthermore, establishing clear, accessible governance guidelines and ethical policies—such as rules around data privacy, source verification, and avoiding bias—is crucial. This ensures employees understand the guardrails for responsible use, building trust in the technology and confidence in the organization's approach.

The second critical component is focused on **continuous learning and fostering new forms of collaboration**. Since effectively utilizing GenAI relies heavily on asking the right questions, organizations must prioritize training in "prompt engineering" and critical thinking to help employees become proficient partners with the AI. This shift requires moving away from traditional, siloed roles toward a model where human intelligence guides and refines AI output. Encourage the creation of interdisciplinary teams—mixing roles like data science, operations, and creative design—to brainstorm and develop innovative GenAI applications. By building this culture of cross-functional learning and shared expertise, the organization ensures that the benefits of GenAI are distributed broadly and that the technology is adopted as an organic part of the workflow, rather than a top-down mandate.

2. Executive Alignment and Governance: Setting the Strategic Compass

The foundation of any successful AI initiative rests on strong leadership and clear governance. This initial phase is about setting the “why” and “how” of AI adoption.

- Define AI Adoption Goals: Move beyond vague aspirations to concrete, measurable objectives. For instance, a public-sector organization might aim for a 15% reduction in administrative workload within the first year, freeing up resources for citizen-facing services. Goals should also encompass qualitative improvements, such as enhancing the equity of service delivery by using AI to identify and mitigate biases in existing processes.
- Appoint AI Leads: Departmental AI leads are the linchpins of execution. These individuals should not only possess a strong understanding of their department’s operations but also a keen interest in technological innovation. A case study from a large government agency revealed that successful AI leads were often mid-level managers with a reputation for creative problem-solving and strong interpersonal skills.

- Mandate Regular AI Training for Leadership: To ensure informed oversight, leadership training is non-negotiable. Programs like the “AI for Executives” course at Berkeley Haas School of Business or Harvard’s Professional Development programs offer strategic, non-technical overviews of AI’s potential and pitfalls. This training equips leaders to ask the right questions and make sound decisions about AI investments.

- Ensure Oversight by Executive Committees: An executive committee, composed of a diverse group of senior leaders, should be responsible for the ethical and strategic oversight of all AI projects. This committee would review and approve high-impact AI use cases, monitor performance against defined goals, and ensure alignment with organizational values and public policy.

3. Workforce Strategy & Team Structure: The Human Element of AI

The success of AI in government is about people more than technology. Building AI fluency across the workforce requires a strategy that prioritizes inclusive learning, peer-to-peer knowledge sharing, and practical experimentation. Think back to the early 2000’s when Web 2.0 began to emerge. Professionals were figuring out what it meant to move from static web pages to bi-directional communication through interactive, user-generated platforms. We saw federal agencies leaning into practical, use-case driven adoption. Early adopters were encouraged to experiment and share what worked, sparking informal communities of practice around tools like wikis, blogs, and internal social platforms. These grassroots efforts helped demystify new technologies and fostered a culture of collaboration. Agencies also created sandbox environments, safe spaces where staff could test tools without operational risk. This approach offers a valuable blueprint for AI adoption today: focus on fluency, not mastery, and empower employees to learn from one another through real-world applications.

Building AI-Ready Teams: Structure & Strategy

To parallel AI and the Web 2.0 era, federal agencies, at that time, experienced meaningful shifts in organizational design that offer valuable lessons for AI adoption today. The rise of collaborative tools like wikis, blogs, and internal social platforms led to flatter hierarchies and more horizontal communication, empowering frontline staff to contribute insights and drive innovation. New roles

emerged, such as digital strategists and community engagement leads, reflecting a more integrated approach to technology and public service. Agencies also began experimenting with agile methods, forming smaller cross-functional teams and embracing iterative development cycles. Innovation became decentralized, with departments piloting tools independently and sharing learnings across the enterprise. These changes were supported by updated governance structures and a renewed emphasis on knowledge sharing. As agencies now navigate AI integration, similar principles apply: empower employees at all levels, foster peer-to-peer learning, and design flexible teams that can adapt to emerging technologies.

Unlike the Web 2.0 wave, which focused on collaboration and innovation but rarely addressed ethical concerns, the AI era begins with a fundamentally different mandate: ethics are front and center. Every conversation about AI in government must start with who is building it, who is impacted, and how bias is actively mitigated. This shift has direct implications for workforce strategy. Agencies must prioritize demographic and cognitive diversity in their teams, not only to ensure fair outcomes but to build public trust.

As agencies move from ethical principles to practical implementation, the first step in building an AI-ready workforce is identifying an AI Champion within each team. These individuals don't need to be technical experts; they need to be trusted peers who can translate AI's value into everyday workflows. By giving them a platform to share use cases and providing access to lightweight training and tools, agencies can foster a culture of peer learning and experimentation. As AI tools begin to enhance individual performance, improving speed, quality, and output, managers will face a growing imbalance between early adopters, such as the AI Champion, and those who lag behind. This isn't just about who knows how to use AI; it's about how AI is amplifying overall job performance. To level the playing field, managers must proactively support lagging employees by identifying barriers, offering targeted coaching, and creating low-risk opportunities to experiment.

At the same time, performance reviews should distinguish between AI fluency and job impact, recognizing that higher output may be tool-assisted. The goal is not to penalize slower adopters, but to ensure access to AI's benefits and to foster a culture where learning and adaptation are valued as much as results. Building an AI-ready workforce doesn't start with complex training programs, it starts

with permission to explore. Agencies should focus on real use cases that solve operational challenges. Create safe spaces to experiment, such as pilot programs, internal hackathons, and sandbox environments, where employees can test tools without fear of failure. Encourage peer-to-peer sharing through informal channels like lunch-and-learns, internal newsletters, and collaborative platforms. These grassroots strategies foster confidence, accelerate learning, and build a culture where AI fluency grows organically.

4. Practical AI Adoption Strategies: From Permission to Proficiency

This section focuses on the "how-to" of getting employees to embrace AI tools in their daily work, particularly at the foundational level.

- Permission to Prompt: Many employees, particularly in the public sector, are hesitant to use new technologies without explicit approval. A formal, documented "permission to prompt" policy, communicated from the top down, can alleviate these concerns and signal that AI is a sanctioned and encouraged tool.
- Focus on High-Pain, Low-Risk Tasks: Initial AI adoption should target tasks that are time-consuming and tedious but have a low risk of negative consequences if performed imperfectly by an AI. Examples include:
    - Email and Communication: Drafting emails, summarizing long threads, and generating meeting agendas.
    - Summarization and Research: Condensing lengthy reports, academic papers, or news articles.
    - Document Drafting: Creating initial drafts of memos, presentations, and standard operating procedures.
    - Administrative Tasks: Scheduling meetings, organizing files, and transcribing audio recordings.
- Build Habits via Real Workflows: One-off demos are often forgotten. The key is to integrate AI into existing workflows. For example, a customer service team could start by using an AI assistant to generate initial drafts of responses to common inquiries, which are then reviewed and personalized by a human agent.

- Embed AI into Daily Tools: The less friction there is in accessing AI, the more likely it is to be used. Integrating AI capabilities directly into familiar tools like Microsoft Outlook, Google Workspace, and customer relationship management (CRM) systems is a powerful strategy for driving adoption.
- Ensure that the AI Interface provides features such as folders and cataloging to ensure the availability of previous interactions.

5. Training and Skills Development: Building a Future-Ready Workforce

To responsibly scale AI across the federal workforce, agencies must invest in training that is practical, inclusive, and aligned with policy mandates like the Sandbox Act. This legislation encourages the creation of protected environments where employees can safely experiment with AI tools, test use cases, and build fluency without fear of operational or compliance risk. Training should begin with foundational AI literacy, including AI 101 and hands-on exposure to Generative AI tools like GPTs, and expand into role-specific learning tracks for legal, procurement, communications, and program management staff. A blended learning approach works best: combining sandbox environments, peer-led demos, microlearning modules, and self-paced courses from platforms like

GSA's AI Training Series and InnovateUS. These formats support different learning styles and allow employees to engage with AI in ways that are relevant to their day-to-day responsibilities. As fluency grows, agencies should begin offering intermediate and advanced training tracks for technical and innovation-focused roles, including topics like prompt engineering, model evaluation, agentic AI, and ethical deployment. Certification pathways can help formalize expertise, while communities of practice sustain learning through peer exchange and shared experimentation. Importantly, training must also address the ethical dimensions of AI, including bias mitigation, data privacy, and transparency, to meet OMB guidance. By embedding training into the broader workforce strategy and aligning it with federal directives, agencies can ensure that AI adoption is not only effective but equitable, secure, and trusted by the public.

6. Security and Data Privacy: Protecting Your AI Ecosystem

At the foundational level, establishing core security principles and basic controls is paramount.

- Establish Clear Data Classification Policies: Define categories for data based on sensitivity (e.g., public, internal, confidential, highly restricted). This helps determine appropriate security measures and access controls for data used by AI models. For instance, personally identifiable information (PII) or protected health information (PHI) should be subject to the strictest controls.
- Implement Guardrails for your AI models: Organizations are very likely to have multiple models as no single model can perform all of the functions required by an organization. Guardrails are an essential feature. They ensure that the model operates safely, ethically, and reliably, within defined boundaries.
- Implement Strict Access Controls: Limit access to AI models, data, and development environments only to authorized personnel on a need-to-know basis. This includes using role-based access control (RBAC) to ensure that employees can only access the data and tools necessary for their specific roles. Where possible, access should be derived from existing ACLs.

- Employee Training and Awareness:
  - Mandatory Data Privacy and Security Training
  - Guidelines for AI Output Review: Train employees to critically review AI-generated content, especially for sensitive information. This includes understanding the potential for data leakage if proprietary or confidential information is inadvertently included in prompts or if AI outputs contain unverified or biased data.
  - Reporting Security Incidents: Establish clear procedures for employees to report any suspected security breaches, data privacy violations, or unusual AI behavior. Utilize existing processes and procedures whenever possible.

7. Risk Management & Incident Response Framework

At this level, the focus is on basic awareness of AI-specific risks and initial steps for incident logging.

- AI-Specific Risk Categories (Awareness): Introduce the concept that AI systems fail differently than traditional software. This includes understanding potential issues like:
  - Model drift - Performance degrades over time without warning.
  - Data poisoning - Malicious or corrupted training data affects outputs.
  - Adversarial attacks - Inputs designed to fool AI systems.
  - Hallucination incidents - AI generates false but convincing information.
  - Bias amplification - AI systems perpetuate or worsen existing biases.
  - Prompt injection - Users manipulate AI through carefully crafted inputs.
- Incident Logging: Implement a simple system for logging any observed AI failures or unexpected behaviors. This forms the basis for future analysis.

Implementation Roadmaps (Foundation Level)

- 30-60-90 Day Quick Wins:
  - 30 Days: Define initial AI adoption goals, appoint departmental AI leads, conduct first leadership AI training session, establish “permission to prompt” policy, identify 3-5 high-pain, low-risk tasks for initial AI use.
  - 60 Days: Conduct initial prompt writing and responsible AI use training for pilot teams, embed AI capabilities in one commonly used tool (e.g., email drafting), implement basic data classification for AI-related data.
  - 90 Days: Gather initial feedback from pilot teams, establish an incident logging system for AI-related issues.
- Resource Requirements: Initial budget for training, dedicated time for AI leads, access to basic AI tools (e.g., enterprise-grade LLMs).
- Success Criteria: High engagement from pilot teams, documented “permission to prompt” policy, at least 2 high-pain, low-risk tasks successfully supported by AI, initial logging of AI incidents.
- Common Pitfalls: Lack of clear leadership buy-in, over-ambitious initial projects, inadequate training, ignoring early signs of AI malfunction.

Intermediate Level (Months 6-18): Expanding Capabilities and Control

Having established the fundamentals, this stage focuses on scaling AI use, enhancing security, and building more robust measurement and incident response capabilities.

1. Practical AI Adoption Strategies: Deepening Integration

- Build Habits via Real Workflows (Advanced): Move beyond initial integrations to embed AI deeply into more complex existing workflows. This might involve using AI for advanced data analysis within specific business units or automating parts of a customer service journey with human oversight.
- Encourage Peer-to-Peer Learning and Best Practice Sharing: Formalize platforms for AI Champions and early adopters to share their successes, challenges, and lessons learned. This could include internal newsletters, brown bag sessions, or a dedicated knowledge base.

2. Training and Skills Development: Building a Future-Ready Workforce

Training at this level becomes more specialized and proactive.

- Sector-Specific Workshops: Conduct workshops focused on the unique challenges and opportunities of AI within specific sectors, such as healthcare, finance, or education.
- Encourage Microcredentialing and Modular Courses: The field of AI is constantly evolving. Platforms like Coursera and edX offer a wide range of short, specialized courses and microcredentials that allow employees to continuously update their skills in areas like prompt engineering, AI ethics, and data analysis.

3. Meeting Executive Orders / Policy Requirements: Ensuring Compliance and Alignment

- Align with Government AI Action Plans: As of early 2025, the U.S. government is in the process of developing a new AI Action Plan. Organizations should stay abreast of such developments and ensure their training and use cases align with national priorities for trustworthy and responsible AI.
- Track Adoption Progress: Regular surveys and departmental audits are crucial for monitoring the progress of AI adoption across the organization. This data can be used to identify areas where adoption is lagging and to tailor interventions accordingly.
- Leverage Public-Sector Pilots: Participating in public-sector pilot programs can be an effective way to test and refine AI solutions in a real-world setting, identify scalable models for broader adoption, and contribute to the development of best practices.

4. Measurement and Metrics: Quantifying the Impact of AI

What gets measured gets managed. A robust measurement framework is essential for demonstrating the value of AI and making data-driven decisions about future investments.

- Key Performance Indicators (KPIs):
  - Daily/Weekly AI Usage Rates: Track the percentage of employees actively using AI tools.
  - Estimated Time Savings: Aim for a quantifiable goal, such as an average of 120 hours saved per employee per year. This can be measured through self-reporting, workflow analysis, or A/B testing.
  - Worker Confidence and Satisfaction: Use regular pulse surveys to gauge employee sentiment towards AI and their confidence in using it effectively.
  - Perceived Relevance of AI: Ask employees to rate the relevance of AI to their specific roles and tasks.
  - Before-and-After Adoption Metrics: For each training cohort, collect baseline data on key metrics before the training and then measure the same metrics again after a set period (e.g., three to six months) to assess the impact of the training.

5. Ensuring AI for All

AI has the potential to increase the accessibility of information and systems.

- AI for Accessibility: AI-powered tools can be a game-changer for accessibility. For example:
  - Neurodiverse Users: AI assistants can help with organization, task management, and communication for individuals with ADHD or autism.
  - Non-Native Speakers: Real-time translation and language simplification tools can improve communication and understanding.

6. Security and Data Privacy: Protecting Your AI Ecosystem

Building on foundational security, this stage incorporates more advanced measures.

- Anonymization and Pseudonymization: Whenever possible, anonymize or pseudonymize sensitive data before it's used for training or by AI models. This reduces the risk of re-identification while still allowing the AI to learn from the data. For example, a healthcare organization might use anonymized patient data to train an AI for diagnostic support, ensuring individual patient identities are protected.
- Data Lineage and Audit Trails: Maintain comprehensive records of data lineage, tracking where data originated, how it was processed, and how it's used by AI models. Implement audit trails to monitor who accesses what data and when, allowing for quick identification of any unauthorized activity.
- Secure Model Development and Deployment Pipelines: Protect the entire AI lifecycle, from data ingestion and model training to deployment and monitoring. This involves using secure coding practices, vulnerability scanning for AI frameworks and libraries, and ensuring secure deployment environments for AI models.
- Adherence to Industry-Specific Regulations: Organizations in regulated sectors (e.g., finance, healthcare) must ensure their AI deployments comply with relevant industry-specific regulations and standards, such as HIPAA for healthcare or PCI DSS for financial data.
    - NIST's new Control Overlays for Securing AI Systems (COSAIS) concept paper strategically extends NIST established SP 800-53 cybersecurity framework into AI security. By creating specific overlays for different AI use cases, NIST offers organizations a practical way to adapt existing security controls rather than implementing entirely new frameworks, making AI security more accessible and implementable. COSAIS is built on top of AI RMF and the upcoming Cybersecurity Framework for AI. By leveraging standards that the industry is already familiar with, it makes it easier for organizations to have a clear and consistent way to protect the confidentiality, integrity and availability of their AI systems.

- The Cloud Security Alliance and OWASP's Artificial Intelligence Security Verification System (AVISS) came out with an Agentic AI Core Vulnerability Scoring System. The principle of the scoring system is to address the rising risk associated with AI Agents.

7. Risk Management & Incident Response Framework

Developing more structured incident response plans and integrating risk assessment.

- Incident Response Playbooks: Develop initial playbooks for common AI incident types, including:
  - Detection protocols - How to identify different types of AI failures (e.g., performance degradation, unusual outputs).
  - Escalation procedures - Who gets called when AI systems misbehave.
  - Rollback strategies - Quick reversion procedures for each AI system type.
  - Communication templates - How to explain AI incidents to stakeholders (internal only).

- Risk Assessment Framework (Initial):
  - Impact scoring - Begin to assess the business consequences of different AI failure modes.
  - Likelihood assessment - Initial estimation of the probability of various AI risks occurring.
  - Risk tolerance definition - Early discussions on acceptable levels of AI-related risk.

Implementation Roadmaps (Intermediate Level):

- 30-60-90 Day Quick Wins:
  - 30 Days: Automate basic monitoring for key AI systems, establish version control for prompts and models, implement simple cost tracking for AI usage.
  - 60 Days: Pilot automated testing pipelines for select AI models, conduct cross-functional workshops for specific AI use cases, deploy advanced security measures for sensitive AI projects (e.g., anonymization).
  - 90 Days: Begin ROI measurement for key AI initiatives, standardize data classification for all AI-related data, develop initial incident response playbooks for common AI failures.
- Resource Requirements: Increased budget for specialized training, dedicated data science/engineering support, investment in monitoring and testing tools, legal/compliance consultation for AI regulations.
- Success Criteria: Documented ROI for at least two intermediate-level AI projects, reduction in reported AI-related issues due to testing, formal incident response playbooks in place.
- Common Pitfalls: Inadequate data quality, failing to integrate AI into core business processes, underestimating the need for continuous training, ignoring early warning signs from monitoring.

Advanced Level (18+ Months): Optimization and Strategic Integration

At this stage, AI is deeply embedded in the organization's operations, with a focus on continuous optimization, predictive capabilities, and strategic portfolio management.

1. Implementation and Scaling Tips: From Pilot to Enterprise-Wide Adoption

Successfully scaling AI from a few pilot projects to an enterprise-wide capability requires careful planning and execution.

- Focus on Interactive, Hands-On Formats: Whether virtual or in-person, training and workshops should be interactive and hands-on, allowing employees to experiment with AI tools in a supported environment.
- Use Real-World Examples: Build trust and demonstrate the relevance of AI by using real-world examples and success stories from within the organization.
- Offer Follow-Up Support and Continuous Feedback Loops: The learning journey does not end with a single training session. Establish ongoing support channels, such as a dedicated Slack or Teams channel, and create a formal process for collecting and responding to employee feedback. This continuous loop of feedback and improvement is the hallmark of a learning organization and the key to sustained AI success.

2. Security and Data Privacy: Protecting Your AI Ecosystem

At the advanced level, security becomes highly proactive and integrated.

- Protect Against Model Inversion and Evasion Attacks: Implement defenses against adversarial attacks. Model inversion attacks aim to reconstruct training data from a deployed model's outputs, potentially exposing sensitive information. Evasion attacks manipulate input data to cause a model to make incorrect predictions. Regular security testing, including adversarial robustness testing, can help identify and mitigate these vulnerabilities.

- Regular Security Audits of AI Systems: Conduct periodic security audits and penetration testing of AI applications and infrastructure. This helps identify and address potential weaknesses before they can be exploited.
- Staying Current with AI-Specific Regulations: The regulatory landscape for AI is rapidly evolving. Designate a team or individual responsible for monitoring new AI-specific laws, guidelines, and executive orders (like the U.S. government's AI Action Plan mentioned earlier) and ensuring the organization's AI practices remain compliant.

3. Risk Management & Incident Response Framework

Full integration of risk management and highly mature incident response capabilities.

- Incident Response Playbooks (Advanced): Refine and automate incident response processes, including:
  - Recovery procedures - Steps to restore normal operations quickly.
  - Post-incident analysis - Formal processes for learning from AI failures to prevent recurrence.
  - Automated Communication Templates - Pre-approved communication for various stakeholder groups (internal and external if required).

- Risk Assessment Framework (Mature):
  - Mitigation strategies - Develop and implement specific actions to reduce identified risks, including technical controls and policy changes.
  - Regular risk reviews - Conduct ongoing, systematic assessments of the AI risk landscape, adapting strategies as new AI capabilities emerge or regulations change.
  - Predictive Monitoring: Implement systems that can anticipate potential AI failures or risks before they occur.

Implementation Roadmaps (Advanced Level):

- 30-60-90 Day Quick Wins:
  - 30 Days: Transition to fully automated AI operations for routine tasks, implement predictive monitoring for critical AI systems.
  - 60 Days: Develop and test self-healing capabilities for select AI models, conduct an organization-wide review of AI governance for continuous optimization.
  - 90 Days: Formalize strategic AI portfolio management, conduct advanced adversarial robustness testing for high-impact AI systems.
- Resource Requirements: Continuous investment in R&D, highly specialized AI security and governance teams, advanced AI platforms and tools, ongoing legal and ethical oversight.
- Success Criteria: AI systems demonstrating self-healing capabilities, demonstrable cost savings and efficiency gains through full automation, established strategic AI portfolio management, zero major unmitigated AI security incidents.
- Common Pitfalls: Becoming complacent with AI governance, failing to adapt to new AI technologies, neglecting ethical considerations as AI becomes more autonomous, under-investing in continuous security and risk management.