GENERAL DYNAMICS
Information Technology

RESEARCH STUDY

AI in Full Bloom

Agency Guide to Scaling from Pilot to Production

With AI's transformative potential now more accessible than ever, this report delves into the strategic planning, technological management, and organizational change required to harness it fully – integrating AI seamlessly into agency operations, ensuring long-term, sustainable success in their missions.

In partnership with

Google Cloud

Foreword

In the wake of generative AI's exponential rise beginning in 2022, we are witnessing the "full bloom" of artificial intelligence across the federal government, marking a pivotal shift in how the technology is leveraged to serve the American public and our national security. This surge in interest is not just about the allure of new technology – it's about the transformative potential of AI to redefine how the government operates, makes decisions, and delivers services.

Accelerated by the White House's executive order, which directs a comprehensive approach to lead in AI innovation while managing its risks, the stage is set for an optimistic journey forward. Yet, fulfilling the vision outlined in the executive order presents challenges, including managing data, ensuring model performance at scale, controlling costs, maintaining safety, ensuring fairness, and achieving user adoption. These areas represent both technical and organizational hurdles that agencies must overcome to successfully implement AI solutions.

This report examines the practical aspects of making AI work at scale, building on the excitement of AI's capabilities. It explores how agencies are currently deploying AI, ensuring its safe and secure operation, and overcoming challenges related to scaling, data diversity, and workforce readiness. We illuminate the path forward, highlighting best practices, innovative solutions, and the critical importance of strategic planning and execution.

The journey of AI from pilot to production is as much about navigating organizational and cultural shifts as it is about technology. The goal is to adopt AI safely and responsibly and to integrate it in ways that drive efficiency, foster innovation, and ultimately enhance mission outcomes.

DAVE VENNERGRUND
Vice President
Artificial Intelligence & Data Insights
GDIT

Executive Summary

The rapid advancement and broadening accessibility of AI capabilities presents an unparalleled opportunity for federal agencies to harness this potential for transformative change. However, transitioning from pilot project to operational success in production remains a significant hurdle, characterized by a need for strategic planning, technological adeptness, and organizational agility.

By focusing on the collective experience of agency peers, this report distills a set of best practices to evolving initiatives into scalable, secure, and widely adopted AI systems.

1
SCALABILITY IS KEY TO AI PROJECT SUCCESS

58% of respondents say pilot projects fail to move into full production because of challenges related to scalability. Solutions architected with an eye towards scalability from the start ensure resilience and adaptability, enabling concepts to become operational successes.

2
DATA DIVERSITY HAS A DUAL IMPACT

Data plays a crucial role in AI success, with 38% of respondents indicating successful projects leverage diverse data formats to gain richer insights. However, managing these varied data formats can also lead to issues, as 33% say their projects fail for the same reason. This underscores the importance of robust data management strategies.

3
BRIDGING THE USER ADOPTION GAP

23% of respondents say their projects stall before production likely due to misalignment between AI solutions and user expectations. Addressing this challenge requires a user-centric design philosophy, ensuring AI solutions are accessible, intuitive, and seamlessly integrated into existing workflows, thereby enhancing their adoption and impact.

4
14-MONTH JOURNEY TO PRODUCTION

The journey from pilot to production takes an impressive average of just 14 months, demonstrating the efficiency and focus of agency AI projects. This rapid timeline highlights the importance of being deliberate and thorough to ensure that AI initiatives are not only quickly implemented but also set up for long-term success.

5
USE CASES DRIVE SUCCESS, NOT JUST TECHNOLOGY

With over half of respondents (52%) leveraging AI for data analytics and reporting, the message is clear: successful AI initiatives are those rooted in clearly defined use cases. This focus ensures AI applications are not just technologically advanced but are also deeply integrated with and responsive to the unique needs and aspirations of the organization.

6
SUCCESS THROUGH PARTNERSHIPS

73% of respondents cited external expertise and partnerships as crucial for ensuring technical success. By drawing on a broad spectrum of knowledge and experience, organizations can navigate the complexities of AI development more effectively, paving the way for groundbreaking solutions that drive value.

AI in Action

Agencies have embraced the power of AI, with a wide array of projects now operational, reflecting the deep-rooted integration of machine learning solutions across government operations. These initiatives predominantly focus on harnessing historical, structured data (data in databases or spreadsheets, where everything is organized into rows and columns), a testament to the longstanding familiarity and expertise within the data science community. This trend underscores the commitment to leveraging AI for enhanced decision making, operational efficiency, and service delivery to the public.

AI Project Success

Projects leveraging predictive analytics, at 71%, and machine learning, at 70%, were the most likely to succeed from pilot to production. These capabilities are most mature in the marketplace and often rely on structured data.

Image analytics projects, at 55%, are less likely to succeed, and generative AI, at 65%, is doing well considering its recency and the extra challenges unstructured data (like emails, videos, images, and social media posts, where content is not organized in a predefined manner) present.

	SUCCEEDED	FAILED
Predictive analytics and forcasting	71%	29%
Machine learning and deep learning	70%	30%
Text analytics and natural language processing	66%	34%
Generative AI and chat	65%	35%
Chat bots and cognitive agents	57%	43%
Image analytics	55%	45%

EXPERIENCE WITH AI

Agencies have varying degrees of experience with AI, ranging from nascent to mature stages of adoption. Dispelling the notion that AI has just emerged, a quarter of all respondents have used AI for more than five years. On average, agencies have been using AI for nearly three years.

10%	40%	26%	25%
Less than 1 year of experience	1-2 years of experience	3-5 years of experience	Over 5 years of experience
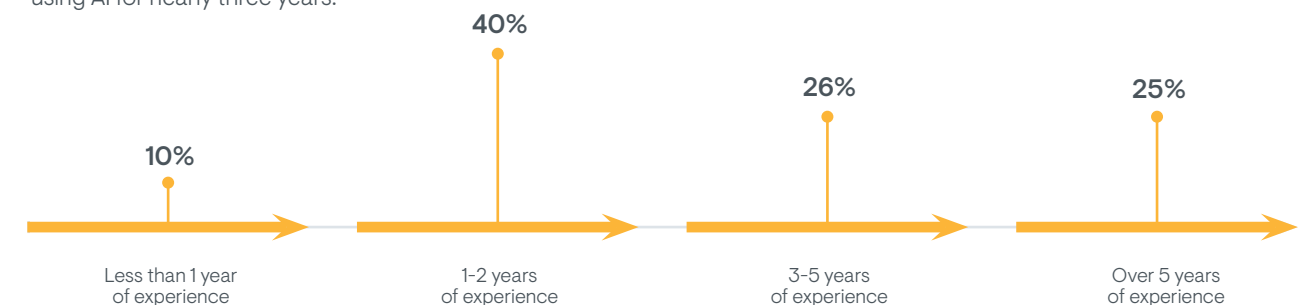
14-Month Journey to Production

While the adoption of AI is in its nascent stages, advancing projects from pilot to production is a deliberate process. Respondents indicated that projects tend to progress rapidly through the pilot phase, w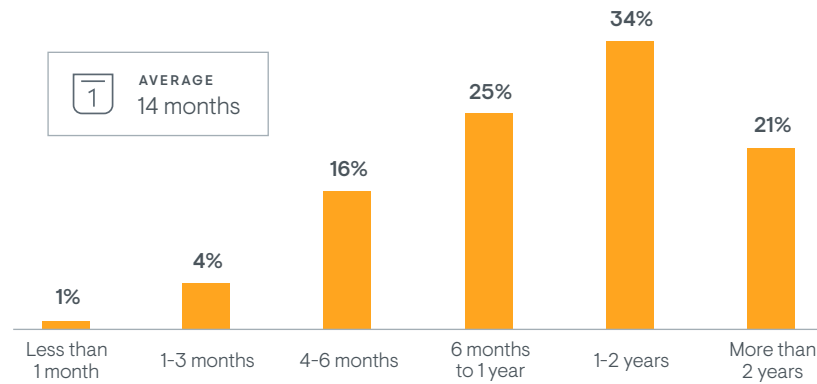ith the majority transitioning to production within two years or less. On average the timeline extends to about 14 months to make the full transition.

Length of Time from Pilot to Production

				34%	
			25%		21%
AVERAGE 14 months		16%			
	4%				
1%					
Less than 1 month	1-3 months	4-6 months	6 months to 1 year	1-2 years	More than 2 years

This rapid progress can be attributed to a variety of factors.

- The scope and complexity of the pilot might generally be narrow, focusing most AI projects on specific, well-defined problems that have limited scope and complexity.
- Heightened focus and diligence in AI projects, due to their unique complexity and risk, might contribute to their accelerated advancement.
- The involvement of specialized data science talent could also be a determining factor, where their expertise is likely to enhance the project's success rate.
- Advanced cyber and DevOps processes are presumed to facilitate safe delivery, ensuring projects are not only swift but secure.

Primary Reasons to Use AI

The reasons for using AI vary widely, but they are largely aiming to enable faster decision making, with 68% ranking it among their top three. This objective is complemented by interests in increasing efficiency and access to real-time data.

	TOP REASON	2ND REASON	3RD REASON
Enable faster decision making and finding insights in data	17%	29%	22%
Increased efficiency and productivity	22%	12%	11%
Access to more robust, updated or real time data	23%	12%	8%
Improved customer and citizen experience	6%	14%	7%
Cost savings	10%	11%	6%

AI USE CASES

The U.S. government's public inventory of AI programs, last updated in September 2023, serves as a valuable extension to our original findings.

There are 710 AI projects documented across 21 unique government departments. The Departments of Energy and Health and Human Services are in the lead.

DEPARTMENT OF ENERGY	DEPARTMENT OF HEALTH AND HUMAN SERVICES	DEPARTMENT OF COMMERCE
178 projects	157 projects	49 projects

These projects have a balanced distribution across different phases of the development lifecycle. (Not all projects have their phase described.)

OPERATION AND MAINTENANCE	DEVELOPMENT AND ACQUISITION	INITIATION
91 projects	81 projects	74 projects

Diverse technological approaches, known as AI "techniques", are being utilized across these projects. We've grouped these techniques into broader categories for a more structured summary. (Not all projects have their approach described.)

MACHINE LEARNING	ARTIFICIAL INTELLIGENCE	DATA ANALYTICS
107 projects	70 projects	52 projects

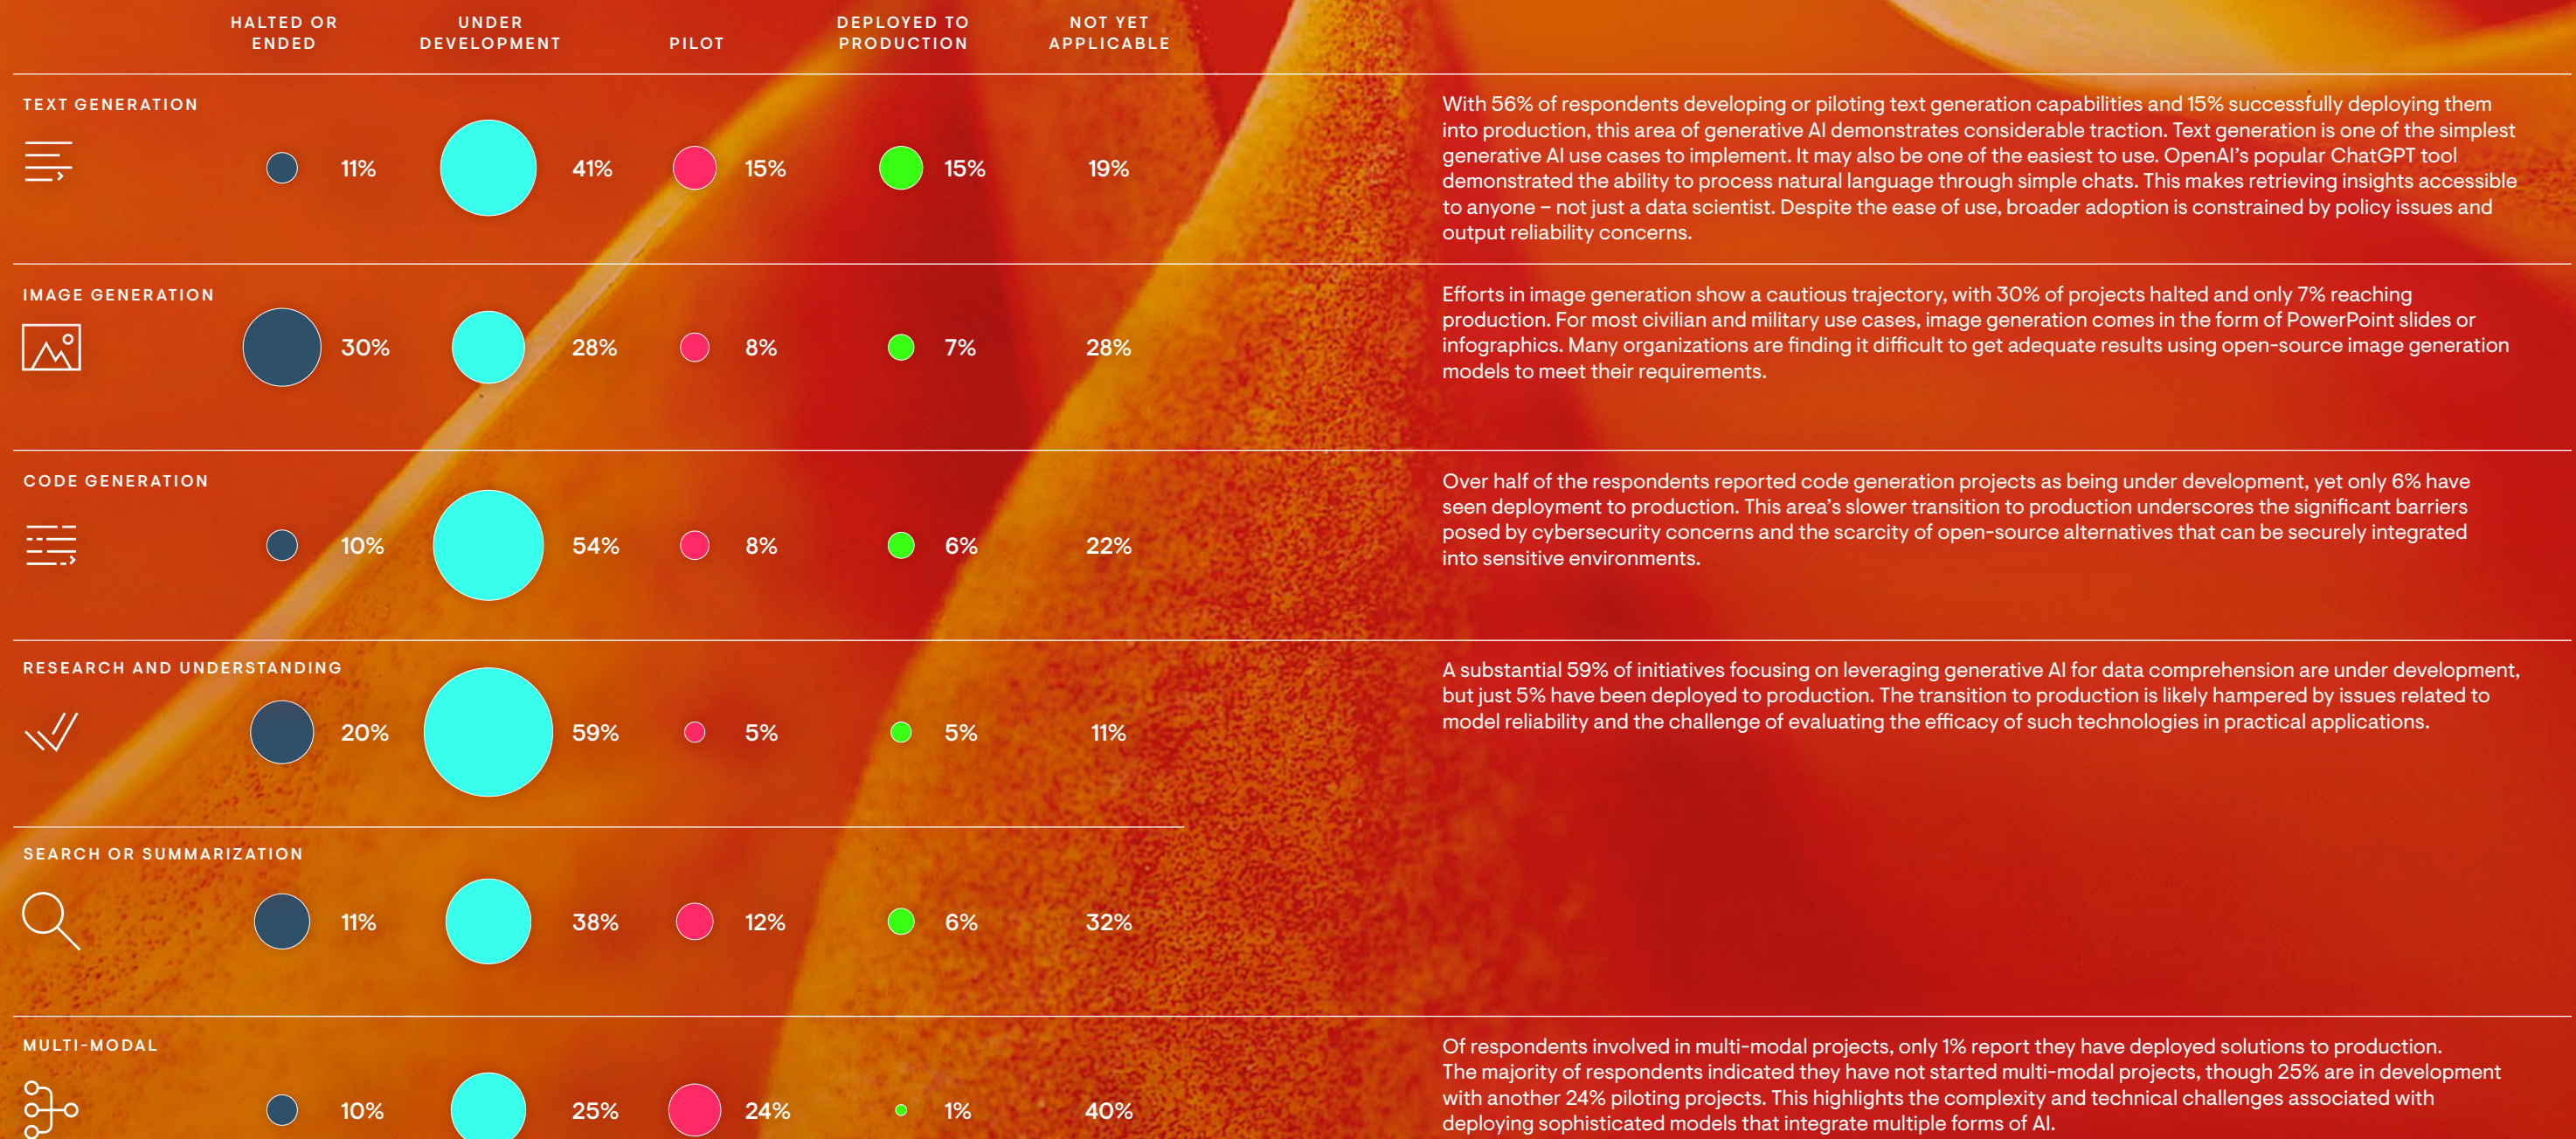| NATURAL LANGUAGE PROCESSING	COMPUTER VISION	ROBOTICS AND AUTOMATION
45 projects	18 projects	1 project

OTHER - 199 PROJECTS

SOURCE: AI.GOV

The Impact of Generative AI

The ascent of generative AI since 2022 has significantly shaped federal agencies' approach to AI, spotlighting both its vast potential and the hurdles to its full-scale deployment.

The data highlight several challenges in adopting generative AI, primarily related to policy alignment. Organizations often await further guidance from higher authorities to ensure their policies are in step with broader governmental directives. In the absence of such policies, many find themselves crafting their own approaches. Nearly three-quarters of respondents reported that their organizations have developed policies grounded in established industry frameworks like the National Institute of Standards and Technology AI Risk Management Framework.

	HALTED OR ENDED	UNDER DEVELOPMENT	PILOT	DEPLOYED TO PRODUCTION	NOT YET APPLICABLE
TEXT GENERATION	11%	41%	15%	15%	19%
IMAGE GENERATION	30%	28%	8%	7%	28%
CODE GENERATION	10%	54%	8%	6%	22%
RESEARCH AND UNDERSTANDING	20%	59%	5%	5%	11%
SEARCH OR SUMMARIZATION	11%	38%	12%	6%	32%
MULTI-MODAL	10%	25%	24%	1%	40%

TEXT GENERATION: With 56% of respondents developing or piloting text generation capabilities and 15% successfully deploying them into production, this area of generative AI demonstrates considerable traction. Text generation is one of the simplest generative AI use cases to implement. It may also be one of the easiest to use. OpenAI's popular ChatGPT tool demonstrated the ability to process natural language through simple chats. This makes retrieving insights accessible to anyone – not just a data scientist. Despite the ease of use, broader adoption is constrained by policy issues and output reliability concerns.

IMAGE GENERATION: Efforts in image generation show a cautious trajectory, with 30% of projects halted and only 7% reaching production. For most civilian and military use cases, image generation comes in the form of PowerPoint slides or infographics. Many organizations are finding it difficult to get adequate results using open-source image generation models to meet their requirements.

CODE GENERATION: Over half of the respondents reported code generation projects as being under development, yet only 6% have seen deployment to production. This area's slower transition to production underscores the significant barriers posed by cybersecurity concerns and the scarcity of open-source alternatives that can be securely integrated into sensitive environments.

RESEARCH AND UNDERSTANDING: A substantial 59% of initiatives focusing on leveraging generative AI for data comprehension are under development, but just 5% have been deployed to production. The transition to production is likely hampered by issues related to model reliability and the challenge of evaluating the efficacy of such technologies in practical applications.

MULTI-MODAL: Of respondents involved in multi-modal projects, only 1% report they have deployed solutions to production. The majority of respondents indicated they have not started multi-modal projects, though 25% are in development with another 24% piloting projects. This highlights the complexity and technical challenges associated with deploying sophisticated models that integrate multiple forms of AI.
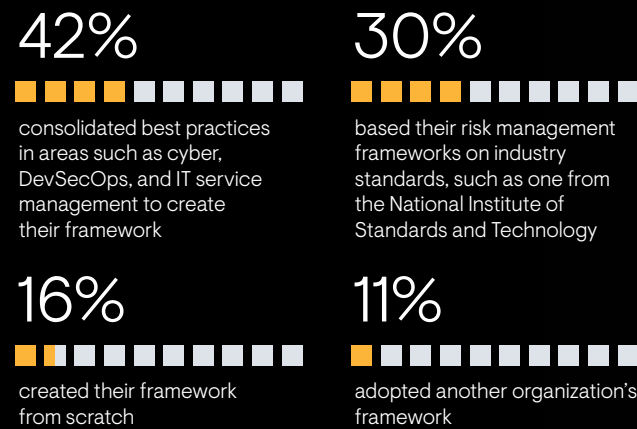
Steering AI with Safety

The Executive Order for the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, issued in October 2023, is a call to federal agencies to accelerate their adoption of AI and do so in a safe manner. It outlines key responsibilities, policy adherence – including NIST's AI Risk Management Framework – and the establishment of AI governance and inventory, complete with specific tasks and deadlines to ensure a framework for secure and ethical AI usage.

The response of agencies to the White House's executive order on AI showcases a concerted effort to align with national goals for leading and responsibly managing the promise and risks of artificial intelligence. This journey, reflective of the executive order's comprehensive strategy, underscores the federal commitment to not just participating in the AI revolution, but leading it with a focus on safety, security, and societal benefit.

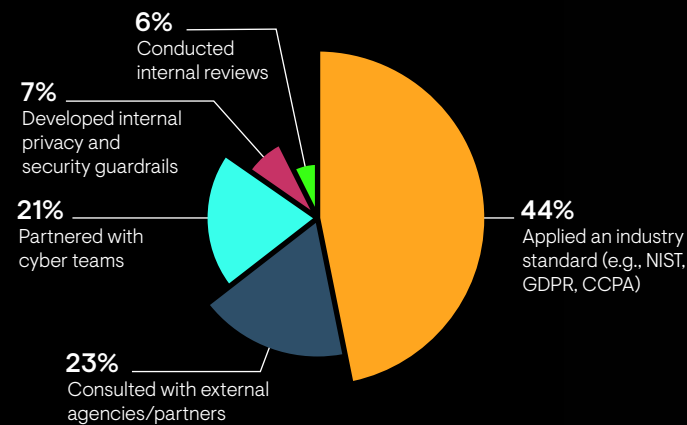### Accuracy, Fairness, and Trustworthiness

44% of respondents identified applying a risk management framework as their primary method to ensure accuracy, fairness, and trustworthiness. Using industry and custom benchmarks was a distant second at 24%.

There are different approaches to risk management frameworks for different people, though.
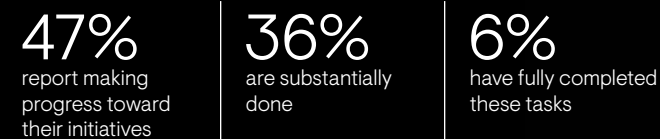
42%
consolidated best practices in areas such as cyber, DevSecOps, and IT service management to create their framework

30%
based their risk management frameworks on industry standards, such as one from the National Institute of Standards and Technology

16%
created their framework from scratch

11%
adopted another organization's framework

Security and Privacy

The executive order's focus on security and privacy has seen agencies adopting industry standards, engaging in collaborations with external agencies or partners, and partnering with cyber teams.

6% Conducted internal reviews

7% Developed internal privacy and security guardrails

21% Partnered with cyber teams

23% Consulted with external agencies/partners

44% Applied an industry standard (e.g., NIST, GDPR, CCPA)

AI Governance

A substantial portion of agencies are strengthening AI governance – this includes pivotal steps like appointing chief AI officers and developing AI use case inventories, aimed at crafting robust frameworks for ethical AI use.

47% report making progress toward their initiatives

36% are substantially done

6% have fully completed these tasks

Tackling AI Scaling Challenges

This report is aptly labeled as a guide to scaling from pilot to production – 58% of respondents reported that the reason their AI solution didn't survive was that it couldn't scale to production demands.

This difficulty suggests that agencies may need to focus on developing robust, scalable AI architectures and strategies from the outset. Considering scalability during the pilot phase itself will help ensure smooth transitions to larger-scale implementations. These considerations include the cost to train and run models, inference volume, users, and data volumes.

Within this all-encompassing journey, there are other reasons projects struggle, from the human to the technical.

58%
of respondents cited struggles due to scalability issues
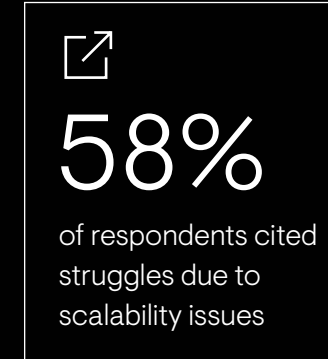
23% Lack of end user adoption

This indicates a potential disconnect between AI solution design and the actual needs or preferences of end users, emphasizing the importance of user-centric design, alignment with operational objectives, and change management strategies.

20% Insufficient funding

Financial constraints hinder the ability to scale AI projects, highlighting the need for clear business cases and ROI to secure adequate budget for AI initiatives, as well as committed stakeholders, and demand from a user community.
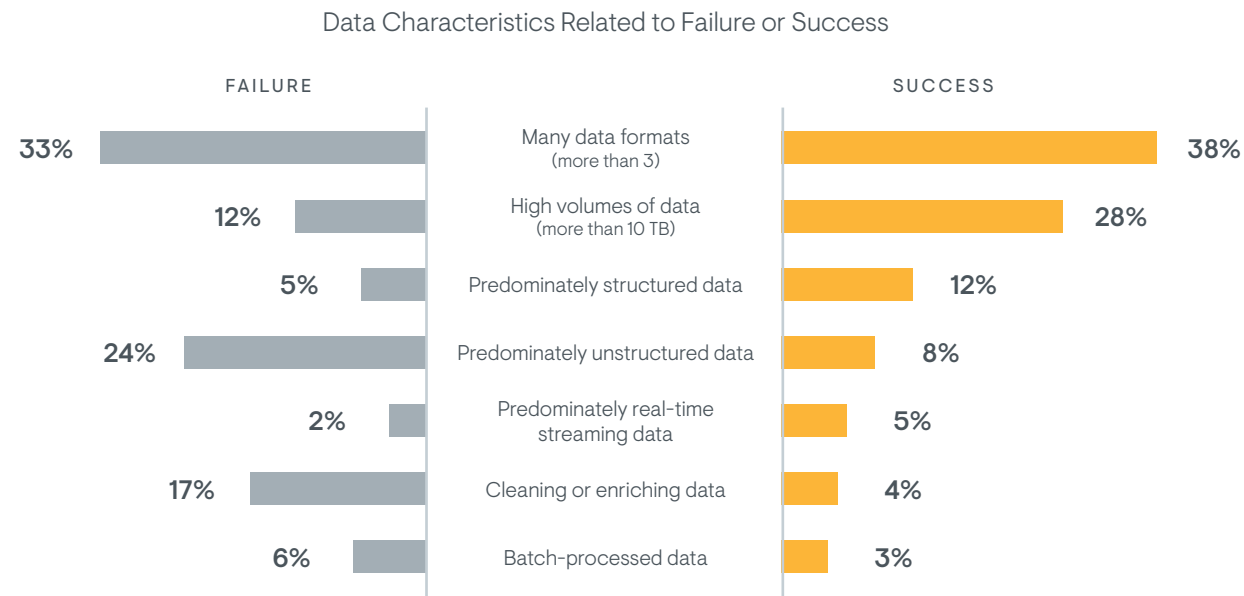
18% Unable to achieve authorization to operate (ATO) from a security point of view

This reflects the critical importance of incorporating security considerations and compliance into AI solutions from early stages to ensure smooth transition to production environments.

A Data Doubletake

Data is the lifeblood of artificial intelligence, providing the raw material from which AI systems learn, evolve, and make decisions, fundamentally shaping the capabilities and effectiveness of AI solutions.

Data Characteristics Related to Failure or Success

FAILURE		SUCCESS
33%	Many data formats (more than 3)	38%
12%	High volumes of data (more than 10 TB)	28%
5%	Predominately structured data	12%
24%	Predominately unstructured data	8%
2%	Predominately real-time streaming data	5%
17%	Cleaning or enriching data	4%
6%	Batch-processed data	3%

Agencies face a delicate balance with AI data: They must harness the richness of diverse data types while managing the complexities they introduce.

(+) On one hand, having a variety of data formats – text, images, video, audio, and structured data, among others – is seen as the leading characteristic that contributes to the success of AI solutions, cited by 38% of respondents. This diversity likely enables more comprehensive AI models, allowing for richer insights and more accurate predictions.

(−) On the other hand, 33% of respondents identify the same diversity in data formats as the primary characteristic leading to failures in AI initiatives, suggesting that that while a variety of formats can also pose significant challenges in integration, processing, and consistency, potentially leading to complexities that hinder effective AI deployment.

Harmonizing Volume and Complexity

Exploring the impact of high volumes of data and the distinction between unstructured and structured data offers further insights into how these elements influence the success and challenges of AI initiatives. High volumes of data are valued or their potential to improve AI solutions, offering a rich foundation for developing more accurate and insightful models.

Historically, AI and data science were roughly divided by data type.

STRUCTURED DATA
Structured data challenges used machine learning techniques to great success. Today, classic machine learning and predictive analytics account for more AI projects in production and are more likely to succeed.

UNSTRUCTURED DATA
Unstructured data are traditionally far more difficult to prepare and resulted in fewer deployments. Recently, deep learning and large language models have evened the playing field, accelerating unstructured data use.

Lessons in AI Deployment

Facing the question, "If your organization had to do it again, what would you do differently?" yielded over 300 free text responses.

We applied AI, natural language processing (NLP), and large language models (LLM) to extract insights from the responses. These insights were revised for clarity by our AI subject matter experts. This approach demonstrates the effectiveness AI-assistance for rapid, scalable analysis of textual information.

WORKFLOW INTEGRATION
Minimize workflow disruptions and integration challenges by leveraging a team with subject matter expertise. Establish clear metrics and feedback loops to optimize AI solutions continuously. Streamline operations and prioritize change management, including user training and support.

"Create a thorough change management plan to reduce disruption."

"Ensuring smooth compatibility by integrating pilot models with our current IT frameworks."

DATA MANAGEMENT
A dedicated team focusing on API integration and data pipelines can help ensure a smooth transition to production. Invest in upfront data quality checks and standardization. Insist on strong security measures to safeguard data.

"Better data preparation."

"Conduct rigorous security assessments to detect and resolve vulnerabilities before deployment."

USABILITY AND STAKEHOLDER ENGAGEMENT
Enhance transparency in training and use case development across workflows. Provide comprehensive training on AI capabilities and procedures. Address AI model explainability to facilitate debugging and troubleshooting, prioritizing methods that are inherently transparent.

"Engage key stakeholders across different departments earlier and more actively."

"More training, more communication about the goals and expected outcomes."

PLANNING AND BUDGETING
The foundation of any successful AI initiative lies in meticulous planning and budgeting. From the broad strokes of project outlines to the granular details of financial allocations, this phase is pivotal. Consider areas including infrastructure planning, performance evaluation strategies, scalability planning, and user training and support ensures successful project execution.

"AI project planning and budgeting from the outset to end."

"As it can be expensive, especially for large-scale projects, it requires meticulous budget planning."

PERFORMANCE EVALUATION
Organizations frequently face significant scaling challenges when moving multiple pilot projects to full production. Integrating pilot models with existing IT frameworks and continuously measuring performance can make AI implementations more practical. Proactively addressing issues during the pilot phase enhances production performance, supported by rigorous quality checkpoints, testing protocols, and validation procedures.
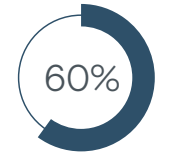
"Choose an architecture that scales efficiently with increasing data volume and computational demands."

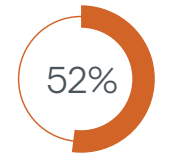"Conduct thorough scalability assessments."

THREE KEYS TO ENSURING TECHNICAL SUCCESS
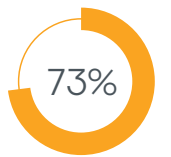
73% Use of expertise or partnerships

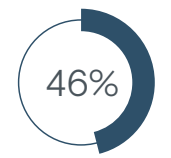**60%** Explicit data governance and data quality requirements

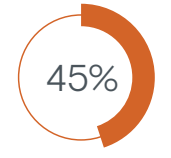**52%** Utilized a trusted commercial end-to-end AI platform

THREE KEYS TO ENSURING AI ADOPTION

53% Integration with existing work flows

46% Support resources/mechanisms for providing user feedback

45% Utilized a trusted commercial end-to-end AI platform

Next Steps

Navigating the transition from pilot projects to a successful full-scale deployment presents a unique set of challenges and opportunities for agencies.

1

ENHANCE AI RISK MANAGEMENT

Organizations should establish comprehensive risk management frameworks that align with both OMB and specific agency guidance to navigate the complexities of AI applications. Utilizing AI inventories and leveraging the expertise of AI offices can provide valuable insights and oversight, ensuring that AI deployments are conducted in a manner that minimizes risk and maximizes benefit. This structured approach to risk management is essential for building trust and confidence in AI solutions across stakeholders.

2

PRIORITIZE DATA GOVERNANCE AND SECURITY

The cornerstone of any successful AI implementation is robust data governance and management. As AI technologies increasingly handle larger datasets, organizations must prioritize scalable management practices and stringent security measures to protect data integrity and privacy. Implementing effective data governance frameworks will ensure that data is accurate, accessible, and secure, facilitating the development of AI solutions that are both powerful and trustworthy.

3

DEFINE USE CASES BEFORE SELECTING AI TECHNIQUES

Traditional AI approaches have a proven track record in production environments, offering reliability and a degree of predictability. Meanwhile, generative AI represents an emerging area with growing application numbers, indicating its potential for innovative solutions. Selecting the AI methodology that is best aligned with the defined use case ensures that the technology application is driven by need and potential impact, rather than novelty.

4

EQUIP THE WORKFORCE FOR AI INTEGRATION

The successful deployment of AI across an organization requires a well-informed and educated workforce. This includes not only educating executives on the value and potential of AI but also ensuring that users understand how to apply and use AI safely and effectively. AI developers, DevOps, and security teams need to be equipped with the knowledge and skills to deploy AI responsibly. Where skill gaps exist, partnerships with external experts can be invaluable. A comprehensive education and training strategy ensures that all levels of the organization are prepared to harness the benefits of AI while mitigating its risks.

5

ENGAGE USERS AND SECURE STAKEHOLDER SUPPORT

Maintaining the interest and support of stakeholders and ensuring continuous engagement with end-users throughout the AI project lifecycle are key to its success. Stakeholder support is essential for securing the necessary funding and resources, while end-user involvement will produce AI solutions that are closely aligned with real-world needs and usability requirements. This dual focus ensures that AI initiatives are both well-supported and effectively tailored to deliver maximum value.

6

SEAMLESSLY INTEGRATE AI INTO EXISTING SYSTEMS

For AI solutions to deliver their intended benefits, they must be seamlessly integrated into existing workflows and infrastructure. This may require careful change management to address workflow modifications and ensure user adoption. Additionally, AI implementations often demand significant data and compute resources, necessitating careful budgeting and planning. Collaborating with system owners, security, and infrastructure teams from the outset is crucial for a smooth integration process, ensuring that AI solutions enhance, rather than disrupt, organizational operations.

GDIT

About GDIT Digital Consulting

General Dynamics Information Technology (GDIT) stands at the nexus of digital consulting and mission-centric innovation in the public sector. GDIT's Digital Consulting Practice supports agencies, across more than 4,000 programs, navigating intricate landscapes and harnessing the transformative potential of AI, cybersecurity, cloud solutions, and emerging technologies. Through a network of GDIT's Emerge Labs, we advance technology research and development by enabling teams to interact, test, and demonstrate the latest technologies, collaborate with over 100 industry partners, and prototype advanced solutions. Partnering with visionary leaders across the public domain, we craft strategies that catalyze digital evolution, drive sustainable modernization, and position our clients at the forefront of excellence.

About GDIT

GDIT is a global technology and professional services company that delivers consulting, technology, and mission services to every major agency across the U.S. government, defense and intelligence community. Our 30,000 experts and consultants extract the power of technology to create immediate value and deliver solutions at the edge of innovation. We operate across 30 countries worldwide, offering leading capabilities in digital modernization, AI, machine learning, cloud, cyber, and application development. Together with our clients, we strive to create a safer, smarter world by harnessing the power of deep expertise and advanced technology. More information about GDIT is available at www.gdit.com.

About Google Cloud

Google Cloud is the new way to the cloud, providing AI, infrastructure, developer, data, security, and collaboration tools built for today and tomorrow. Google Cloud offers a powerful, fully integrated and optimized AI stack with its own planet-scale infrastructure, custom-built chips, generative AI models and development platform, as well as AI-powered applications, to help organizations transform. Customers in more than 200 countries and territories turn to Google Cloud as their trusted technology partner. More information about Google Cloud is available at cloud.google.com/gov.

Contact

GENERAL INQUIRIES

Christopher Aiello
Senior Marketing Manager
GDIT
Email: christopher.aiello@gdit.com

MEDIA INQUIRIES

Jay Srinivasan
Senior Public Relations Manager
GDIT
Email: jayendran.srinivasan@gdit.com