Put AI to work for federal application modernization with IBM and AWS

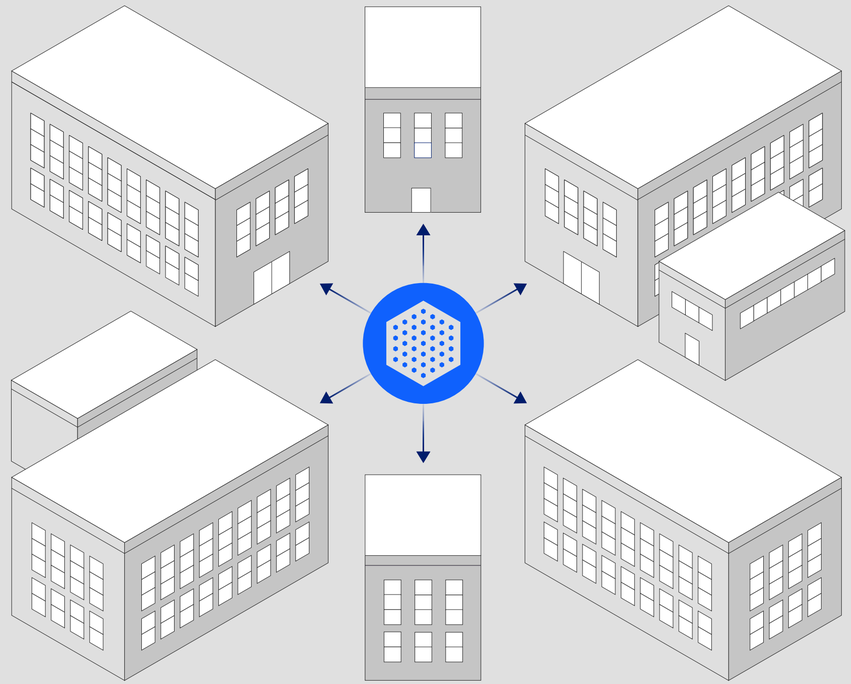

Contents

01

Application modernization in the age of AI

The human urge to modernize. It’s what has driven us for centuries to build new worlds—roads, bridges, cities and financial systems—and invent new technologies. It drives us today to update and optimize our digital infrastructure. Modernization offers the promise of achieving more through improved productivity.

For federal agencies, that promise can mean more agility, more efficiency in their digital infrastructure and more productivity from their developers—and a greater chance of mission success.

Through application modernization with a hybrid infrastructure, which combines on-premises and cloud solutions, federal agencies also have the potential to:

- Update the capabilities of existing applications to meet current technology standards
- Achieve faster resolution of issues
- Improve citizen services
- Create new value
- Build trust

These examples are some reasons why application modernization is a key enterprise imperative for federal agencies.

But the promise of modernization isn't without its challenges. In a rush to meet urgent needs, some federal agencies are now facing siloed cloud or muilticloud strategies that have led to increasingly complex IT environments. Due to this kind of hybrid-by-default approach, the modernization process can be slow and seldom in lockstep with mission imperatives and available resources. To date, only a small number of federal agencies have modernized their essential workflows, related applications, systems and access to data across their organizations.

Barriers to application modernization can be many, including:

- Technical debt accumulated over years of budget shortfalls
- The need for diverse skills and experience escalated by siloed computing platforms
- Cybersecurity and data privacy considerations
- Increasingly complex IT landscape
- Aging applications

What’s hindering application modernization projects?

Why application modernization projects fail, according to respondents of a study conducted by IBM Institute for Business Value:[1]

High costs	57%	of surveyed leaders say they are challenged by financial factors
Technical debt	51%	say they’re challenged by technical factors
Skills gap	45%	say they’re challenged by expertise factors

Fortunately, AI provides new capabilities to help ease these challenges. From increasing developer productivity to accelerating application modernization and operational optimizations across infrastructures, federal agencies can address many strategic outcomes in modernization projects by adopting an AI-first strategy.

Federal CIOs are tasked with enabling their agencies' missions, including reaching agency modernization goals and protecting their infrastructure from rising cyberthreats. If you're one of these IT leaders, application modernization isn't just an IT buzzword—it's an imperative that can help you achieve your goals. How can federal leaders address these challenges to succeed in their modernization efforts at a higher rate?

Two words: generative AI.

To drive transformational success and take full advantage of generative AI (gen AI), leaders must first focus on addressing their IT weaknesses by adopting a hybrid-by-design approach to their hybrid architecture. This strategy pivots from determining business limitations due to technology to purposefully building an architecture that's ready for innovation. An intentional multicloud approach, better aligned to business outcomes, can help supercharge modernization efforts and maximize ROI.

Let's explore why generative AI is key to application modernization and how CIOs can implement it with a hybrid-by-design approach, to simplify and accelerate their modernization journeys.

02

Putting gen AI to work for application modernization

While the benefits of AI are too often mythologized, there's no denying that it's one of the most revolutionary technologies of our time. It's at once transformative and disruptive, holding the potential to both unlock immense economic value and bring about innovative improvements in our lives and work.

The introduction of gen AI has driven up that potential—and exponentially so. Many federal agencies are already working with or planning to adopt gen AI. With wide-ranging capabilities that span different business processes and operations, not to mention the breadth of potential new services, it's clear why many agencies are readily embracing gen AI.

But how can your federal agency, with its unique mission, constraints and operating parameters, successfully apply gen AI?

Success lies in choosing the right foundation models—enterprise-grade, domain-specific foundation models over consumer-focused general-purpose models. A model—whether it's a large language model (LLM), IT automation model or one of the many other options—are typically trained on a wide variety of agency-owned data and tuned to perform mission-specific tasks.

With enterprise-grade AI, federal agencies can bring these models to life to implement new capabilities, such as semantic search, code generation and knowledge management. These capabilities are developed to not only enhance automation and increase efficiency in different processes, but also help augment your staff's skills and knowledge.
To run those models smoothly, you'll want to ensure you're building an intentional hybrid architecture, which allows you to both modernize and place applications effectively on the best infrastructure—private or public—to optimize the application portfolio. Gen AI, with its capacity for automating tasks and generating code, seamlessly integrates into this infrastructure, enabling greater productivity for teams.

There's one caveat to keep in mind while using gen AI to extend developer skills. Just as authors oftentimes need to edit content after using autocorrect tools, developers will still need to refine their work and the work of gen AI. As always, they will make the call on whether the output—code, text, image, etc.—should be accepted as is, modified or completely rejected. However, gen AI can help them fast-track certain steps in the development process.

Scaling AI across the agency safely

Despite the widespread enthusiasm for using gen AI, many federal agency leaders still have concerns with adoption of the technology.

At IBM, we see a few common concerns about AI adoption within federal agencies:

- Data lineage or provenance
- Data security
- Data privacy

How your agency succeeds with gen AI is influenced by the process used to select, govern, analyze and apply data across the organization. Huge volumes of data from different sources are used to train gen AI models, so implementing governance, management and ethical frameworks that operate end-to-end is critical. These frameworks are key to safely and responsibly adopting AI that is supportive of mission outcomes while remaining incredibly secured and in compliance with federal regulations.

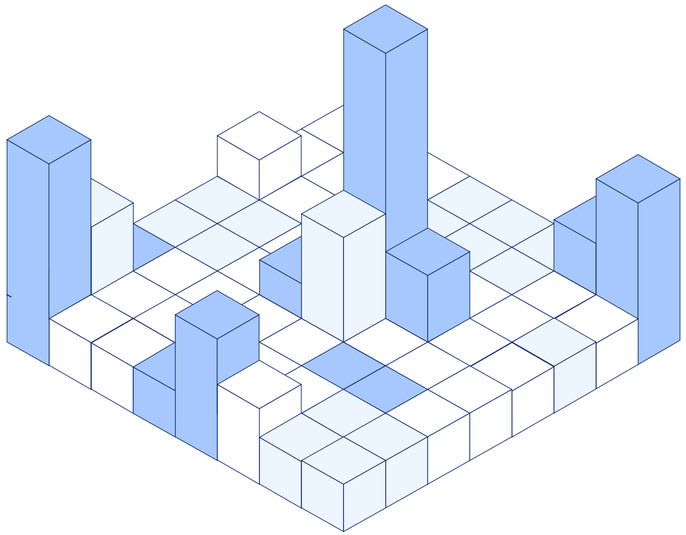

Scaling AI process

As gen AI evolves, federal agencies and public sector organizations should seek to acknowledge and address the ethical considerations, data privacy issues and potential for data misuse by implementing strong governance frameworks and controls.

A hybrid model can also enhance security measures by enabling organizations to establish consistent security practices across IT estates—including core, edge and cloud environments. Businesses can implement tailored security protocols and compliance standards, ensuring data integrity and regulatory adherence throughout the modernization process. When used responsibly, this technology can revolutionize the way federal agencies and public sector organizations deliver on their missions, helping them make data-driven decisions faster, streamline processes and enhance the experience of those they serve.

IBM has long followed core principles grounded in commitments to trust, transparency and fairness to guide how we handle any of our client's data and insights—and how we develop and deploy new technologies.

To continue this practice in the age of AI, IBM has developed a multidisciplinary, multidimensional approach, that embeds ethical principles into AI applications and processes. With IBM's Principles of Trust and Transparency and Artificial Intelligence Pillars of Trust as the foundation for our AI ethics initiatives, we're helping people and organizations adopt AI responsibly. These initiatives are accomplished with clear purpose, and in alignment with the Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.

Many government organizations are in the early stages of regulating gen AI, and it is crucial to remain in compliance with anticipated regulations. Creating a foundation built on organization-wide practices that ensure high levels of security, privacy and anti-discrimination should be the focus. Additionally, federal agencies should consider any regulatory, security or compliance requirements associated with various data sources, especially when handling sensitive information such as personally identifiable information (PII) or healthcare data.

Federal agency IT teams should work closely with their legal teams to ensure non-technical requirements, such as data lineage, are considered as well. United States federal laws, such as HIPAA for healthcare, FERPA for education and FedRAMP for federal information, dictate specific requirements that must be met in the US. For any application it should be assumed that gen AI needs to adhere to these rules, as would any other project.

IBM guiding principles for AI ethics

Principles for Trust and Transparency

The purpose of AI is to augment human intelligence.
We believe that AI should make all of us better at our jobs, and that the benefits of the AI era should touch the many, not just the elite few.

Data and insights belong to their creator.
Clients' data is their data, and their insights are their insights. We believe that government data policies should be fair and equitable, prioritizing openness.

Technology must be transparent and explainable.
Companies must be clear about who trains their AI systems, what data is used and what goes into their algorithms' recommendations.

Pillars of Trust

Explainability
Good design does not sacrifice transparency in creating a seamless experience.

Fairness
When properly calibrated, AI can assist humans in making fairer choices.

Robustness
As systems are employed to make crucial decisions, AI must be secured and robust.

Transparency
Transparency reinforces trust, and the best way to promote this practice is through disclosure.

Privacy
AI systems must prioritize and safeguard consumers' privacy and data rights.

What does gen AI mean for traditional AI?

In the past, organizations have approached AI as an add-on, with the end goal being digital transformation and cloud modernization.

Now, AI is becoming the centerpiece of transformation. As a data point in the commercial space, 75% of business leaders surveyed believe competitive advantage will depend on who has the most advanced gen AI.[2] But harnessing the potential of AI to fundamentally transform application modernization requires a mix of vision and technology. In the case of federal agencies, they too need to put AI to work at the strategic core of their strategy—not just add it on to existing systems—to solve enterprise challenges and help achieve their mission objectives. It's time to move from applications +AI to AI+ applications.

So, does this shift in technology mean there's a need to replace traditional AI solutions with the latest gen AI for application modernization? IBM AI leaders say no.

Traditional AI models that use conventional machine learning (ML) and rules-based models have different capabilities and serve different functions including, but not limited to:

Predictive analysis

Security and compliance

Automation

Gen AI models, on the other hand, use foundation models to autonomously generate content based on the data they were trained on. Gen AI and traditional AI are, in fact, complementary technologies—use them together to help accelerate and achieve your agency's modernization goals.

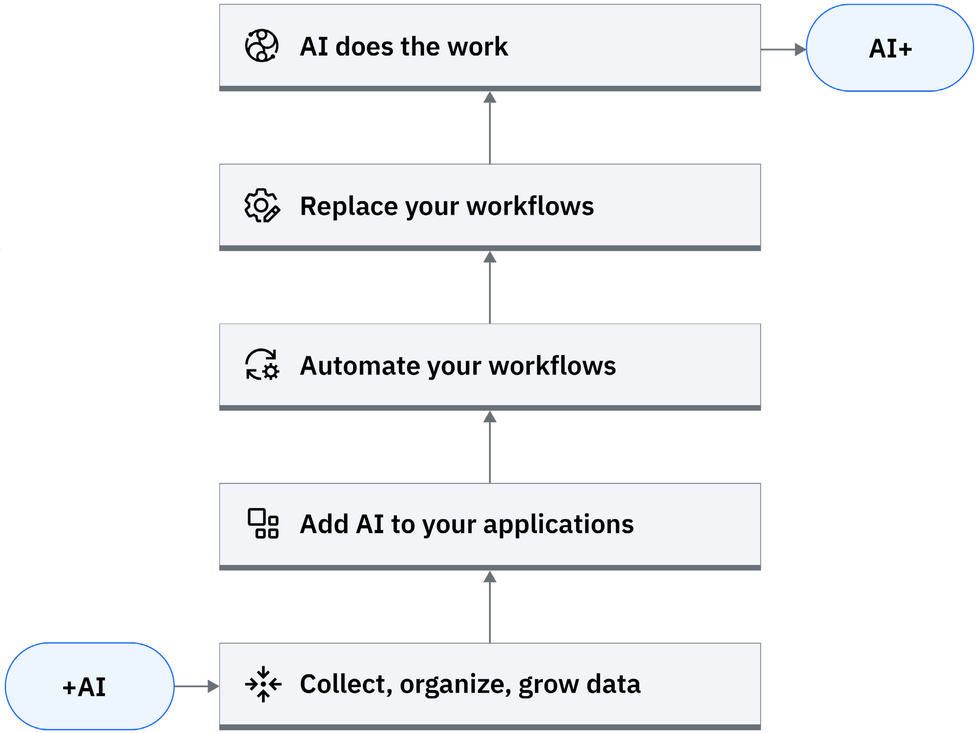

Traditional AI
Applies predefined rules and algorithms to specific sets of data to solve problems, make predictions and automate tasks. Potential application modernization tasks using traditional AI include:

- Capacity planning
- Resource utilization
- Cost optimization
- Performance monitoring

Generative AI
Generates new content based on patterns learned from the data it was trained on. Potential application modernization tasks using generative AI include:

- Code generation
- Best practice recommendation
- Context-sensitive automation
- Code debugging

79% of responding executives say using generative AI in app modernization projects will increase business agility.[3]

Many federal CIOs already recognize the obvious merit in adopting gen AI to progress their many mandates. Now is the time to embrace a new perspective and learn why gen AI could be the solution to many modernization challenges.

Standing up existing systems with modern applications can compound the complexities in the IT environment. Using gen AI instead to tackle modernization issues can help simplify and accelerate the entire application modernization journey. Gen AI-powered code conversion, code generation, code reverse engineering and transformation planning can help address common modernization challenges, improve employee productivity and reduce costs.

Tackling technical debt with generative AI

Technical debt can accumulate across various aspects of an agency's IT infrastructure, including code, architecture and documentation. Identifying and resolving the complexity that comes with this debt can become an intimidating task.

Gen AI offers a solution. When you use it to resolve technical debt, you can accomplish tasks and achieve outcomes that were otherwise impossible. Enterprise-grade AI provides capabilities that improve issue classification, generate code for issue resolution and set up context-sensitive automation—helping to significantly cut down the time developers spend fixing code.

With AI-enabled automation, developers can also reduce the time they spend provisioning cloud infrastructure, applying patches and performing maintenance. From implementing faster code debugging and enabling better document generation to driving automation efforts, gen AI is designed to help reduce technical debt and accelerate modernization efforts one prompt at a time.

Bridging the skills gap with generative AI
The shortage of in-house traditional technical skills is a real challenge many federal CIOs must address today. But with technology evolving faster than ever, it's equally important that IT staff acquire modern skills—including ones related to gen AI and the new disciplines it's creating—such as prompt engineering.

The solution to this problem may seem simple enough: reskill, automate and adapt. But is it still simple if it costs you time? Make use of the code generation capabilities of generative AI and provide your teams a new, faster way to work with a variety of skills and languages.

With these capabilities, teams are positioned to translate code from one language to another or create original code, even when they don't know the specifics or syntax of a language. They also can easily summarize code—typically existing code, which may have been poorly documented—and provide the necessary context without any assistance from the original developers who may no longer be accessible. They can also write automation playbooks with AI-generated recommendations—all with the security and data privacy federal agencies need.

Gen AI in action

The IBM CIO Organization, as “client zero,” used IBM watsonx™ Code Assistant for Red Hat® Ansible® Lightspeed in a tech preview and experienced firsthand the benefits of putting AI to work for application modernization, including:[4]

50% reduction in Ansible Playbook development effort

60% of Ansible Playbook content automatically generated

10x expected increase in the number of people who can help produce Ansible Playbooks going forward

IBM Consulting® also used the technical preview of IBM watsonx Code Assistant for Red Hat Ansible Lightspeed and improved developer productivity with generative AI:

30% reduction in Ansible Playbook development effort while maintaining quality, compliance and resiliency

How to use AI for application modernization

As with any complex IT transformation project, application modernization also happens in phases, including critical workflows that present fitting opportunities for applying gen AI.

Modernization phase	**What happens?**	**Which key workflows align with these phases?**
Advisory	You assess your application and ecosystem landscape and create the modernization roadmap.	Transformation planning
Planning	You develop the project plan and set the modernization milestones.	Transformation planning and code reverse engineering
Transformation	You modernize the applications and migrate them to newer platforms and architectures.	Code generation and code conversion
Testing	You validate the modernized applications to ensure they are fully functional and meet the performance and security requirements of your agency.	Code generation and code conversion

What are workflows?

Workflows are processes and tasks that occur in a sequential order, from start to finish, through which people and enterprises accomplish their work. The workflows discussed in this section are instances where gen AI can be infused into an application modernization process to standardize and optimize outputs, enriching what teams can achieve.

Teaming with IBM Consulting gives you access to services that can help your organization prepare for its transformation journey, meeting you where you are—whether it's creating a custom strategy, modernizing business-critical applications or developing and managing new applications.

Let's look at the 4 main gen AI workflows for application modernization.

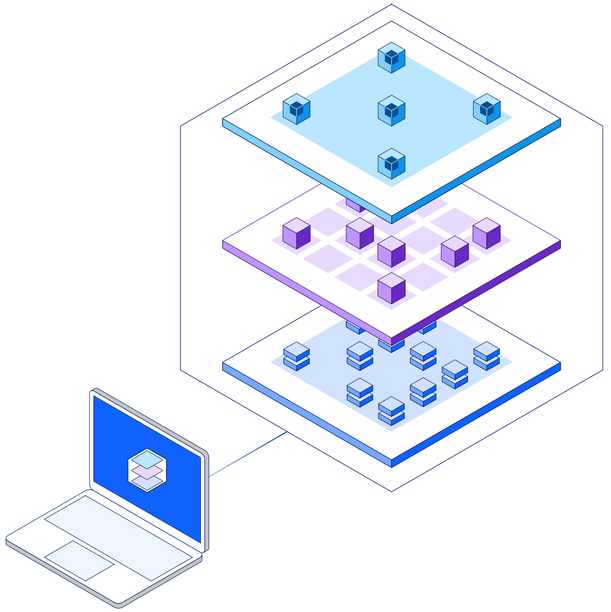

01

Transformation planning

Determining the transformation strategy, including ROI estimation, should take place during the advisory and planning phases. This is when the modernization approach is defined based on both business and technical aspects of applications. Within the advisory and planning phases, organizations define a modernization strategy based on the business and technical properties of the applications and calculate expected ROI. Today, organizations may rely on extensive rules encoded in spreadsheets and other tools to plan. But when there are hundreds of servers and applications to modernize, how do you manage the complexity?

Throughout the industry, technical leaders are considering how gen AI can be adopted to radically streamline transformation planning.

Gen AI may help:

- Accelerate discovery and disposition by processing hundreds of documents and summarizing them if required.
- Create modernization roadmaps with detailed transformation plans and target architecture designs to achieve a hybrid-by-design ecosystem.
- Generate reference architecture, including diagrams and documentation that can be converted into Terraform code to create a landing zone for the transformed applications.
- Help to understand the overall application to make informed modernization decisions.

Amazon has created several gen AI solutions including Amazon SageMaker, Amazon Bedrock and Amazon Q.

Amazon Bedrock is fully managed service offering a selection of foundation models from leading AI companies, including AI21 Labs, Anthropic, Cohere, Meta, Stability AI and Amazon. These offerings are available within a single API that also boasts a broad set of capabilities to build generative AI applications with security, privacy and responsible AI.

Amazon Q is a new gen AI assistant built specifically for use in a workplace setting. This tool can use an organization’s existing data, systems and expertise to provide employees with fast, tailored answers and insights to help them work more efficiently.

Trained on 17 years of Amazon Web Services (AWS) knowledge and experience, Amazon Q can help developers build AWS-based solutions faster. It can also offer guidance on optimal AWS services for a use case, explain code, suggest optimizations, diagnose issues and accelerate upgrades. Amazon Q is also built to:

- Connect to repositories and systems to understand an organization.
- Answer questions using authorized data.
- Generate content grounded to provided sources.
- Help individual employees with summarizing information, generating content, structuring meetings and completing tasks.
- Streamline communications and eliminate repetitive work.

02

Code reverse engineering
In the planning phase, you perform reverse code engineering to analyze the business logic that's embedded in the existing code. It helps you understand the codebase, its architecture and any dependencies.

With gen AI, technology experts are exploring ways to summarize the code and analyze existing applications to generate phased modernization plans based on code dependencies and historical project sequencing.

Gen AI can help organizations:

- Help extract business rules to analyze and summarize code
- Help uncover domain models to analyze code and identify domain-specific entities
- Support monolith-to-microservices transformation to analyze monolithic applications and generate microservice recommendations
- Help optimize application refactoring and containerization to identify refactoring opportunities, generate refactored code and suggest containerization configurations

03

Code generation
You focus on code generation in the transformation and testing phases. But what if you are plagued by developer team fatigue, too many repetitive manual tasks, poor code quality and constant customization challenges?

Gen AI can let you create code very quickly, enabling you to accelerate code development with just natural language prompts.

How Gen AI helps

- Write cloud-native code in Java, Quarkus, Spring Boot and other languages.
- Optimize and document COBOL code.
- Build UI code, including back end for front end (BFF) patterns, by analyzing images to generate code compatible with the UI design.
- Generate infrastructure code, including Terraform infrastructure code.
- Create code for configuring container platforms, such as the Red Hat® OpenShift® Platform.
- Develop serverless code for serverless functions using Knative.

IBM offers enterprise-grade code generation capabilities through IBM watsonx Code Assistant—a gen AI-assisted solution purpose-built to help you accelerate development with multiple highly relevant uses cases. Your developers will still play the valuable role of refining their code when using both IBM watsonx Code Assistant for Z and IBM watsonx Code Assistant for Red Hat Ansible Ligthtspeed. While gen AI can help accelerate the development process, the final approval of the code rests in the hands of human experts.

AI-generated recommendations for Mainframe Application Modernization

IBM watsonx Code Assistant for Z helps equip your developers, regardless of skill levels, with the ability to write code through AI-generated recommendations for mainframe application modernization. With IBM Watson Code Assistant for Z, you can help reduce the burden of understanding applications with huge technical depth by generating code documentation with a purpose-built AI assistant, thus helping to manually write more modular and optimized COBOL business services.

For overall IT automation, start with IBM watsonx Code Assistant for Red Hat Ansible Lightspeed and create Ansible Playbooks with automation content recommendations.

04

Code conversion
In the transformation and testing phases, you can also engage in code conversion to translate or migrate code from one format to another. Maybe you have existing applications that aren't compatible with new platforms, but you still need to retain them. But what if your developers don't have the bandwidth or the expertise required for code conversion from existing applications?

Gen AI can help you automate this conversion process and enable it for many languages. For example, convert code from COBOL to Java by generating optimized Java code that matches the functionality of the original COBOL code.

For code conversion, start with IBM watsonx Code Assistant for Z—an AI-assisted mainframe application modernization solution—and enable your developers to incrementally modernize COBOL business services and selectively convert them to high-quality Java code while automatic equivalence testing. It's enabled by a 20-billion-parameter LLM for code, which is on target to become one of the largest generative AI foundation models for code automation.[5]

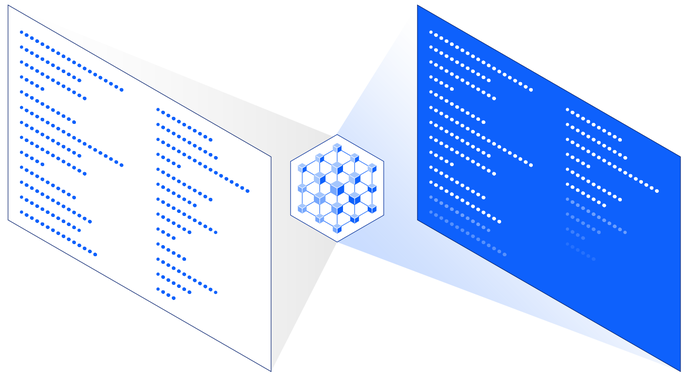

Code conversion

03

Take the next steps

Applying gen AI to application modernization can accelerate your agency’s transformation journey and set in motion a new era of efficiency that enhances mission execution.

But scaling gen AI across the agency is still a challenge for many federal CIOs. And doing it with a focus on trust, compliance and security is an even bigger challenge.

Here are 7 recommendations for federal CIOs to successfully adopt gen AI for application modernization.

01

Ensure the right infrastructure is set up intentionally, with the necessary security, privacy and resiliency controls to support the right AI use case for your agency.

An effective AI journey starts with an intentional hybrid approach. Choose an open, hybrid cloud platform empowered by Red Hat OpenShift Platform that enables deployment across different locations and offers a consumption-based model for both on-prem infrastructure and cloud-delivered services.

Do not over invest or over provision infrastructure for AI. Instead, align the right infrastructure to the right AI task at hand. Make sure the infrastructure meets your organization’s standards for security, privacy and resiliency.

Make sure the infrastructure meets federal and agency standards for security, privacy and resiliency. By integrating AI capabilities within hybrid environments, organizations can enhance their application modernization efforts, ensuring efficient data management and robust security protocols.

IBM’s IT infrastructure solutions—from servers and mainframes to storage systems and software—have been designed to help support the next generation of AI technology and follow a clear set of build guidelines.

02

Start with a clear roadmap and develop a well-articulated plan before you implement the modernization and migration journeys.

Understand how the applications are aligned to your agency and discover any dependencies to ensure your transformation strategy doesn't disrupt normal agency operations.

IBM Consulting can help you discover your entire landscape—from the platform, operating systems and middleware you use to the different applications you run—to define your AI roadmap and tailor it to your transformation journey.

03

Evaluate the right AI tools.

As federal organizations weigh adopting AI tools, they identify the core areas where gen AI can transform operations and services to fulfill stakeholder needs. Understand how organizations can responsibly create value by selecting the use cases and tools that align with both their mission goals and compliance—and constituent obligations placed on them.

04

Pursue both the easy wins and opportunities that were previously off limits, such as applications and processes in core systems.

Seize the low-risk, high-visibility opportunities. Applying gen AI to already-modernized applications can help you more readily highlight its benefits. Direct your efforts toward core systems where gen AI can deliver bigger, more strategic results. Focus on past initiatives that represented the most mission value but were abandoned due to cost or complexity.

IBM Consulting can help you rethink your strategy and modernize with purpose. With a focus on AI and hybrid cloud, we can implement and scale advancements to reinvent your agency's workflows.

05

Seed new teams with experienced talent and also track and measure the impacts of gen AI on developer productivity.

Extend gen AI modernization to opportunities across the agency. Close the gap between IT and operations and form new partnerships that drive technology modernization and mission performance, irrespective of roles.

Measure and reward fast cycles of modernization made possible by gen AI. Use this data to demonstrate the value of gen AI to other leaders and drive consistent and accurate automation adoption across your agency.

IBM watsonx Code Assistant for Red Hat Ansible Lightspeed can help novice developers to write Ansible Playbooks and automate tasks while removing the burden of low-level task creation from more experienced automators.

06

Get your data ready for modernization projects.

Evaluate what specific knowledge your AI system requires. Understand whether the data your application requires is available for consumption and mitigate biases in the data used to train AI models.

To make the most effective use of gen AI, you need a clear understanding of your organization's data sources and desired outputs. Define the collection of knowledge your AI system requires, such as specific data points for a chatbot, event summaries or content alignment with learning standards.

07

Take a strategic approach to AI ethics, ensuring that platform policies are based on the principles of transparency, trust and fairness.

Establish, publish and enforce agency-wide policies, practices and guidelines for federal employees using the technology.

Instill a culture of ethics and AI governance to enable holistic, end-to-end management of AI solutions.

IBM can help you deploy the right AI technology from IBM and AWS in the right business area, designed for maximum impact, with governance and security guardrails built into the process.

04

Get started with IBM and AWS

A partnership built on AI Expertise

IBM and AWS can help you strategize, modernize, build and manage applications with the power of AI and hybrid cloud, creating an intentional approach.

The IBM and AWS partnership, together with Red Hat, brings a unique combination of leading enterprise AI, cloud, infrastructure and open-source technologies can help agencies scale AI workloads quickly and responsibly.

The IBM Consulting Center of Excellence (CoE) for generative AI has more than 21,000 IBM data and AI consultants with a strong track record in AI who are already collaborating with thousands of global clients and partners to shape the future of AI.

Our deep AI expertise is complemented by over 24,000 active AWS certifications, which are helping mutual clients to deploy generative AI. IBM is a Premier Tier Partner for AWS and has 19 AWS competencies—including Government Consulting—and 17 validated service delivery programs (SDDs) including AWS GovCloud (US) Delivery.

IBM has built experience and expertise in Amazon's gen AI technologies, including Amazon SageMaker, Amazon Q and Amazon Bedrock.

With Amazon SageMaker JumpStart, an ML hub equips organizations with access to algorithms, models and ML solutions with a high degree control over how clients can discover, explore, an deploy opensource foundation models such as Llama, Falcon, GPD-J, Hugging Face and beyond.

Now, it's time to propel your federal agency to a new level of innovation through the power of AI.

IBM and AWS bring both technology and the expertise to help you get started with generative AI for application modernization.

Register for a discovery session with IBM and AWS AI experts and learn to:

- Build generative AI applications with enterprise-grade security and privacy with industry-leading foundation models using IBM wastsonx, Amazon SageMaker JumpStart or Amazon Bedrock.
- Customize foundation models using your organization's data to deliver differentiated experiences.
- Train models and run inference and get the most price-performant infrastructure for gen AI using a broad set of GPUs from AWS.

Get started →

1. Modernizing applications on hybrid cloud, IBM Institute for Business Value | Research Insights, 4 May 2023

2. The CEO’s guide to generative AI: Platforms, data and governance, IBM Institute for Business Value, 12 September 2023.

3. The CEO’s guide to generative AI: Application modernization, IBM Institute for Business Value, 15 August 2023.

4. Using generative AI reduced playbook development effort by over 50% in IBM CIO Organization Pilot, IBM CIO Organization, October 2023.

5. IBM Unveils watsonx Generative AI Capabilities to Accelerate Mainframe Application Modernization, IBM Newsroom, 22 August 2023.

IBM Corporation New Orchard Road Armonk, NY 10504

Produced in the United States of America April 2024