GitLab

Securing our future: Modernizing government DevSecOps with Artificial Intelligence

By Liz Burrows

Table of Contents

Executive Summary

Artificial intelligence (AI) presents remarkable opportunities for enhancing efficiency and fostering innovation. However, challenges related to outdated legacy systems and legislative gaps continue to hinder modernization efforts. These factors underscore the critical need for integrating AI within DevSecOps to meet current demands and future-proof government operations.

Current state

Government agencies are increasingly exploring AI adoption. With **over 400 million people** accessing federal services annually, there's a growing demand for automation and collaboration, as agencies strive to provide exceptional services to the public. AI integration is pivotal for improving operational efficiency, yet agencies face hurdles such as outdated procurement laws, a shortage of skilled professionals, and the challenge of bridging the gap between modern technology and legacy systems. These challenges highlight the need for evolving policies to support AI's role in government DevSecOps practices.

Future vision

Looking ahead to 2025 and beyond, AI is expected to become integral to government operations, particularly in streamlining processes and modernizing legacy code. AI-driven automation will allow agencies to move their focus from mundane tasks to strategic innovation, addressing inefficiencies and managing increased workloads. Anticipated policy changes will align AI advancements with ethical standards and data privacy, ensuring responsible AI adoption. Additionally, AI will enhance security measures by proactively identifying vulnerabilities and streamlining incident responses, driving a new era of government efficiency and smarter cyber security.

What comes next?

AI is already automating routine tasks, enabling DevSecOps engineers to dedicate more time to strategic initiatives, innovation, and the continuous improvement of software development and operations processes. As AI evolves, it will become a key partner for DevSecOps developers, augmenting their capabilities and enhancing productivity. This will drive significant improvements in government efficiency and security. Through thoughtful adoption and integration, AI can transform public sector service delivery, creating a secure, efficient, and innovative government that better meets the needs of citizens.

Introduction

The use of artificial intelligence (AI) in the Federal Government presents a tremendous opportunity for modernizing agency operations and enhancing the delivery of government services to the public. Historically, federal agencies have struggled with modernization, relying on a patchwork of systems that were never truly designed to work in harmony. As the world becomes increasingly interconnected, the pace of innovation in AI, automation, and cybersecurity has amplified pressures on governmental infrastructures, often leaving them ill-equipped to meet the rising demands of the public.

As a result of the heightened expectations coming out of COVID, citizens now expect government services to operate with the efficiency and sophistication of the private sector. This shift in expectations underscores the need for AI adoption and integration within software development to drive automation, enhance collaboration, and meet these demands. However, significant challenges persist, including outdated procurement laws and a shortage of technically skilled professionals.

This whitepaper offers a deep dive into the current and future role of AI in government DevSecOps, focusing on its evolution in critical areas like vulnerability management, incident response, and legacy code modernization. Our vision for the future of AI in government is based on GitLab's many conversations with government customers, takeaways from industry events, and the data from GitLab's annual DevSecOps Survey.

Presently, AI plays a crucial role in enhancing efficiency and driving innovation across the public sector. As we look ahead to 2025 and beyond, AI is poised to advance these applications, further embedding itself in government operations.

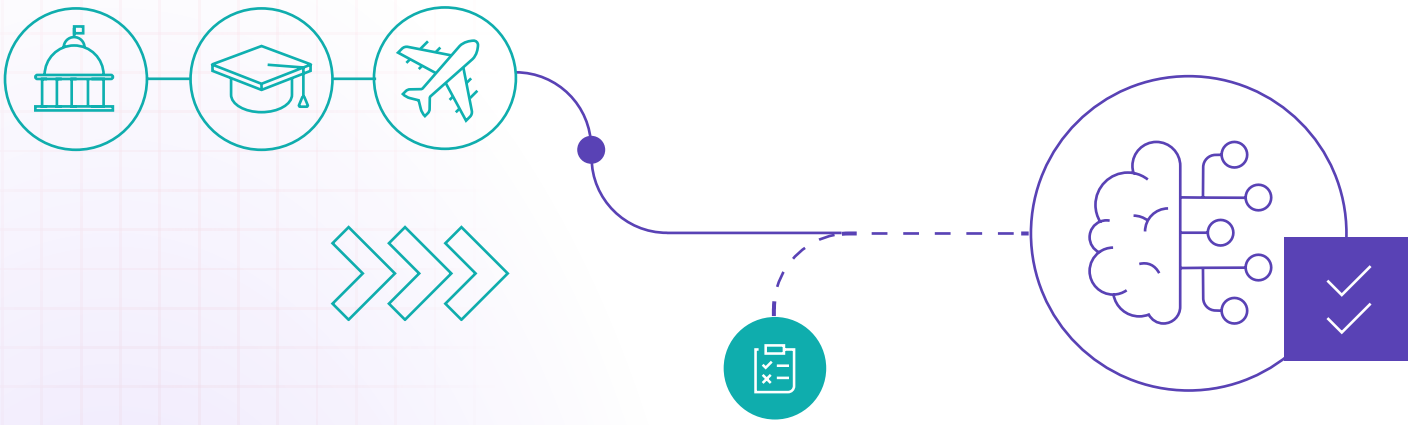

Developer experience

Current state

Government agencies are increasingly exploring AI adoption in DevSecOps to help improve efficiency and the developer experience. According to GitLab's 2024 DevSecOps Survey, 25% of developers in the public sector consider AI assistants crucial for improving their satisfaction. These advancements are critical for teams struggling with heavy workloads, employee retention, and recruiting talent. However, integrating AI into existing systems presents significant challenges, requiring the coordination of various tools and often disrupting established workflows.

25%

of developers in the public sector consider AI assistants crucial for improving their satisfaction

Future vision

As AI-led automation takes on routine tasks, operations within government agencies will become more streamlined, freeing up talent to engage in more strategic endeavors and fostering a focus on creativity. The release of **OMB memo M-24-18**, "Advancing the Responsible Acquisition of AI in Government," directs agencies to improve their responsible acquisition of AI. It provides new guidance for establishing cross-functional and interagency collaboration to manage AI responsibilities, risk, and performance while promoting a competitive AI market. Anticipated policy and governance evolution will see governments crafting adaptive frameworks that balance rapid AI advancements with ethical considerations and data privacy. Without careful integration of AI across the software development lifecycle, government leaders risk creating new silos which could lead to a preference towards platforms that incorporate AI throughout the entire software development lifecycle. As leaders across the software supply chain work to reduce silos and establish a high-trust culture, AI can help in that effort by translating unfamiliar code on the fly and facilitating code review to streamline the review process for better handoffs. This approach, similar to Google Translate, effectively bridges language barriers, enabling smoother communication across the software supply chain.

Modernizing legacy code

Current state

AI is gradually advancing the development of memory-safe applications, fundamentally changing how government leaders address vulnerabilities by transitioning their legacy code to memory-safe code. Many vital applications still run on mainframe systems crafted by individuals who have long since retired, making it challenging to transition these systems to modern architectures. Successful modernization demands that developers adeptly understand legacy code and integrate contemporary solutions into existing frameworks. Yet, integrating AI with older systems poses complexity and risks, potentially introducing vulnerabilities if not managed carefully and often requiring third-party solutions. Furthermore, procurement laws, anchored in outdated strategies, prioritize long-term planning over acquiring technologies aligned with more immediate modernization demands.

Future Vision

In 2025, AI is anticipated to become a cornerstone in government DevSecOps, particularly for refactoring legacy code, enabling more efficient code maintenance and IT modernization initiatives. As AI takes on the heavy lifting of interactive refactoring, developers can efficiently transition applications from languages like Legacy C to RUST while approving real-time code transformations. AI's customizable approach will empower developers to explore multiple solutions, selecting the best fit for their needs. To support this advancement, governments are expected to establish more robust AI regulatory policies and frameworks, ensuring compliance and security. The release of OMB memos M-24-18 and M-24-10 highlights the ongoing efforts to enhance AI acquisition practices. These initiatives will drive cross-functional collaboration and refined management practices, ensuring AI is effectively integrated into agency operations in a more timely manner to meet mission objectives and manage associated risks.

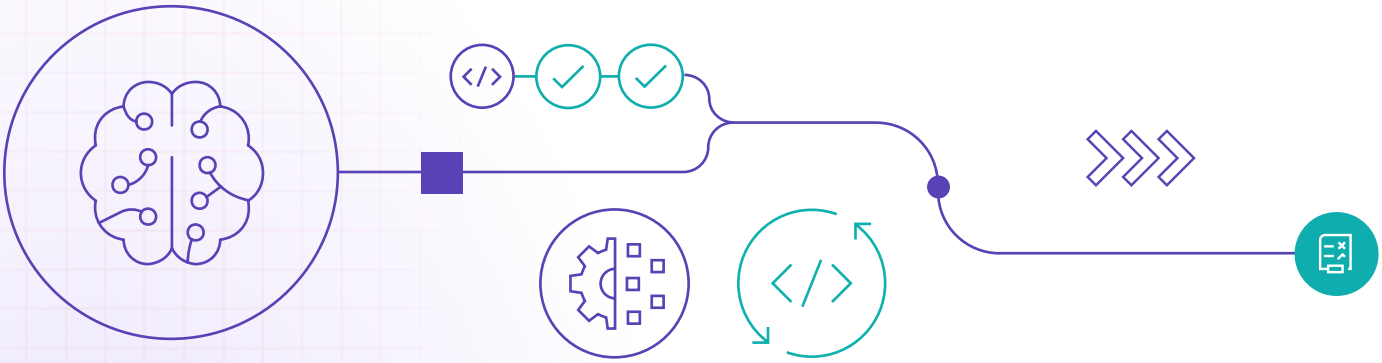

Vulnerability management

Current State

Today, AI-driven vulnerability management can be a double-edged sword for the public sector. While AI can manage numerous tasks, the overwhelming volume of data it processes can inundate security teams, necessitating sophisticated tools to effectively prioritize threats and streamline incident response. While multi-cloud environments and third party dependencies offer increased scalability and performance, these leave agencies vulnerable to potential security issues, demanding vigilant oversight and robust management strategies to mitigate risks.

Future Vision

In 2025, autonomous AI will come into play, proactively identifying and remediating vulnerabilities before code is deployed. This proactive approach will enable AI to dynamically fix potential security breaches before they can be exploited. Enhanced machine learning algorithms will drive advanced threat detection, allowing real-time identification and response to sophisticated cyber threats. Intelligent task execution will enable AI agents capable of understanding and executing complex tasks across projects. As generative AI and autonomous technologies evolve, the **Cyber Incident Reporting for Critical Infrastructure Act** (CIRCIA) released by the Cybersecurity and Infrastructure Security Agency (CISA) will necessitate accelerated vulnerability disclosures and remediation efforts. With over 316,000 entities **projected** to submit an estimated 210,525 CIRCIA reports over the next decade, the urgency to swiftly manage vulnerabilities will be more crucial than ever to prevent exploitation by malicious actors.

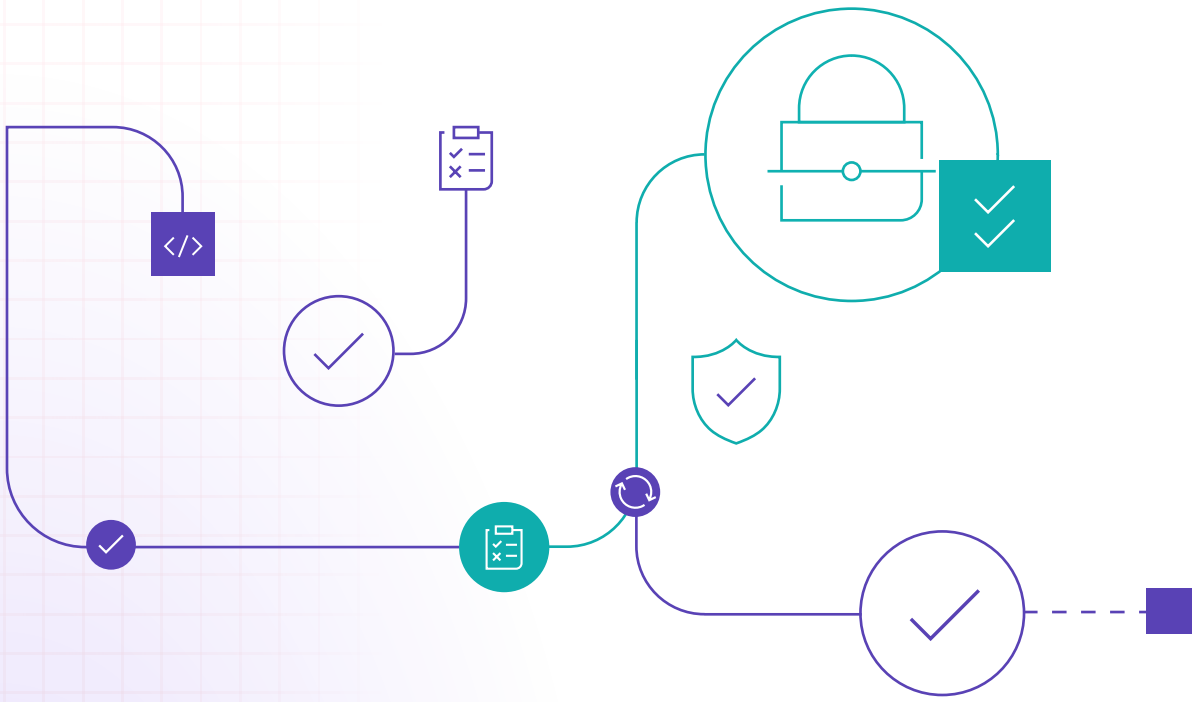

Measuring the impact of AI

Current State

Many organizations are moving beyond AI testing and rolling out large-scale implementations, reaping notable benefits. However, measuring the efficiencies driven by AI remains complex, making the ROI for AI a key focus area for containing excessive spending. AI adoption is contributing to a spike in cloud costs for most enterprises, according to a report by Tangoe. The findings reveal that enterprise cloud costs have surged by an average of 30% in the last year, driven largely by spending on AI applications, with generative AI singled out as a significant factor. Nearly three-quarters of respondents labeled these cloud expenses as "unmanageable." The difficulty in measuring AI efficiencies, especially across various teams and business functions, underscores the need for a strategic approach to AI implementation. This strategy should emphasize a thorough ROI analysis to maximize AI's potential in boosting operational efficiency and driving innovation.

Future Vision

In 2025, organizations will be able to measure ROI more effectively by zeroing in on metrics like AI-generated code acceptance and speed to deploy. By measuring efficiency and productivity gains in these targeted areas, organizations will be able to more accurately assess ROI and make a compelling case for further investment in AI technologies.

AI transparency and regulation

Current State

AI systems present a complex landscape, often criticized for their “black box” characteristics that challenge trust and reliability. This ongoing issue makes troubleshooting difficult and can undermine confidence in AI decisions. Efforts are underway to bolster security, privacy, and accountability. These include initiatives like registering AI models and their data to ensure validation and verification. As data retention policies evolve, they aim to clearly define the models in use and the terms that govern them. There’s a growing expectation that external systems interacting with AI models will align with the rigorous standards set internally. And yet the US government is still grappling with **what adequate regulation for AI looks like**. Notably, however, many organizations are adopting **AI governance frameworks**, akin to those used in cybersecurity and enterprise resource planning, to manage AI risks. These frameworks offer a structured approach to risk assessment and mitigation, ensuring accountability, transparency, and a culture of responsible innovation.

Future Vision

In 2025, governments globally are expected to establish stringent requirements for AI system traceability, mandating comprehensive reports on AI model training data, algorithms, and decision-making processes. This rigorous approach will aim to track the origins and operations of AI systems, thereby fostering greater trust and reliability. Additionally, there is likely to be a concerted push towards developing international standards and fostering collaborative efforts among nations, creating a unified framework for AI transparency. In tandem, organizations will select providers based on their transparency for data usage and the assurances they provide for data protection. There will be greater prioritization around monitoring and control, allowing stakeholders to effectively trace the lineage and modifications of AI models. As part of this evolution, legislation is anticipated to sharpen its focus around issues such as model bias, model learning, and data transparency, ensuring that AI systems are equitable and their data sources are clearly understood and beneficial.

Incident response management

Current State

AI-driven response management in government is mission critical amidst growing cybersecurity challenges. According to GitLab's DevSecOps Global Survey, 67% of security professionals in the public sector find it challenging to get development teams to prioritize vulnerability remediation. Moreover, 60% of these vulnerabilities are discovered after code merges into test environments. Related to these findings, CISA has taken proactive steps by issuing guidance documents that promote AI benefits while ensuring protection against cybersecurity threats. Their roadmap for AI adoption envisions integrating AI into cyber defense missions and provides enhanced guidance for generative technologies. To further these initiatives, CISA has conducted AI-driven security incident response drills, aiming to improve preparedness and refine the application of AI in managing cyber incidents.

Future Vision

AI will accelerate incident response management in government via autonomous actions, such as the detection, analysis, and containment of security incidents. This development will drastically cut down response times, thus minimizing the impact on government operations. Moreover, as cyber incident reports continue to proliferate, AI will play an essential role in automating their processing and review, allowing agencies to manage extensive data volumes more efficiently and ensure timely responses. The exciting prospect of this dual capability is that, by the time an incident report is generated, vulnerabilities will already have been remediated, paving the way for dynamic environment fixes and enhancing overall security readiness.

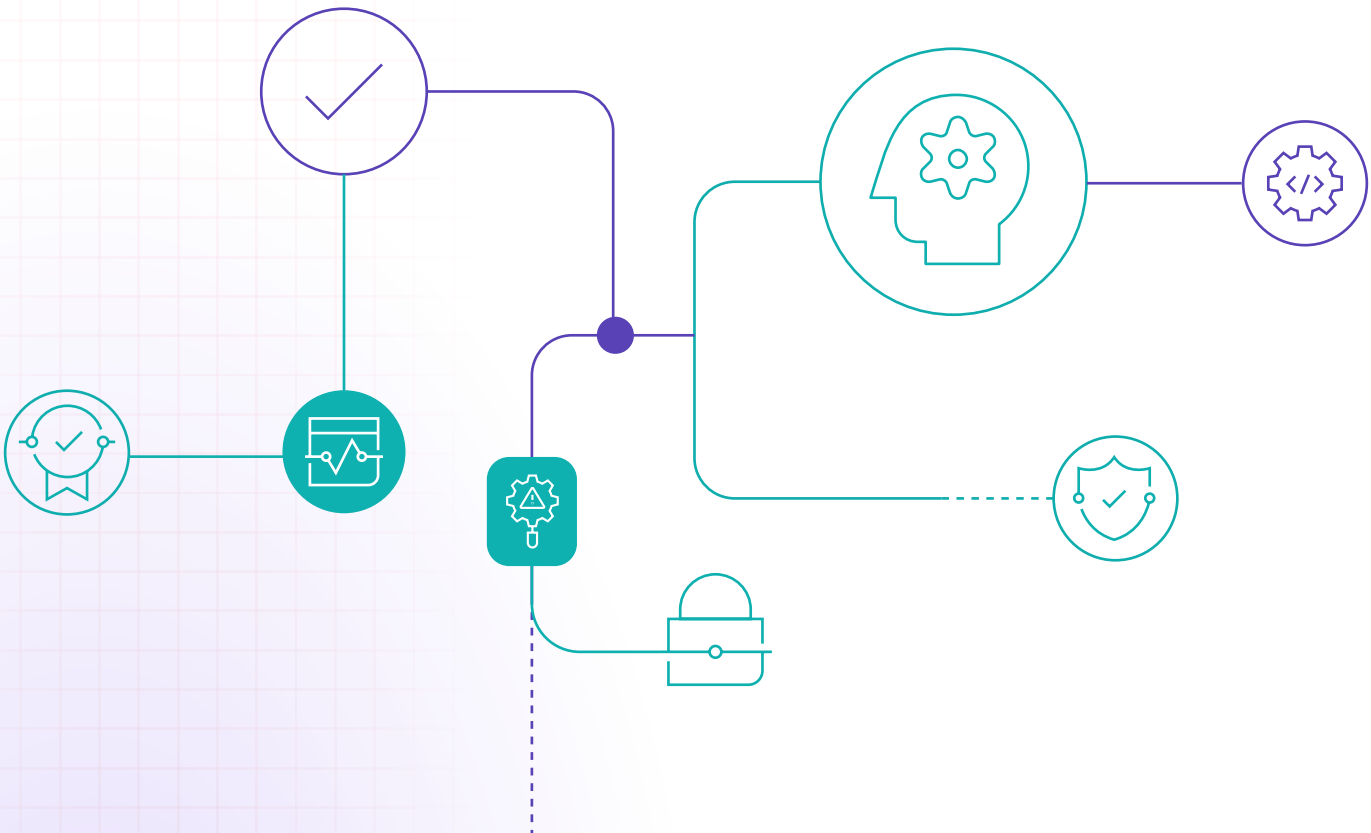

Self-hosted models

Current State

Security is a primary focus for IT teams globally, with the public sector facing even more intense scrutiny than the private sector. While a data breach in a company might lead to negative media attention, a similar incident in a federal agency could result in far more serious repercussions, including appearances before Congress or the leak of national secrets. IT professionals are particularly concerned about potential cyberattacks on cloud providers, compliance with security protocols, and undisclosed security breaches. Due to a lack of AI regulation, agencies are forced to lean into public-private partnerships to navigate cybersecurity and policy issues. This regulatory gap underscores the need for agencies to deepen their understanding of the data without model regulation in place. Self-managed on-premise models address this need by providing enhanced control over data security and ensuring compliance with regulations. They allow agencies the flexibility to leverage AI capabilities without the necessity of migrating to the cloud.

Future Vision

Self-hosted models will provide agencies with increased control over data and processes, allowing for the customization of AI solutions tailored to specific security needs and operational requirements. Driven by the demand for secure, efficient, and customizable solutions that adhere to governmental standards, these models will also meet compliance and privacy needs by operating completely offline within a self-managed environment. They will ensure that all AI network requests remain local, enabling entirely private AI usage. Agencies can configure different local models for various features, empowering government leaders to manage local LLMs that power AI features while keeping all requests within their enterprise network. Separately, the integration of self-hosted AI models with emerging technologies is trending, aiming to automate and enhance the detection of vulnerabilities and the management of security protocols, thereby providing a robust security framework. These advancements are aligned with the 2025 National Defense Authorization Act, which promotes AI modernization and supports innovation to advance national interests.

Conclusion

AI has the power to significantly enhance decision-making, efficiency, and innovation, presenting extraordinary opportunities for governments to deliver on their missions. As AI evolves, it will transition from gradual adoption to becoming a valuable partner in DevSecOps, tackling mundane tasks and configurations so that developers can focus on more creative and engaging development work.This transformation is set to revolutionize the industry by upholding stringent security standards while achieving exceptional efficiency. Guidelines set by NIST and CISA will be top of mind as leaders strive to balance security, speed, and compliance in the face of modernization demands. In its initial phases, AI requires careful monitoring and verification by IT leaders to ensure accuracy and reliability. This scrutiny will help establish a track record of dependability, fostering broader acceptance and reliance on AI technologies. Over time, as these systems demonstrate consistent reliability, organizations can transition towards more autonomous operations, reducing the need for task-based monitoring and human initiated security processes. By strategically adopting AI, government leaders can create a secure, efficient, and cutting-edge DevSecOps environment that addresses the challenges of the digital age.

Let us help

GitLab can help you bring AI into your organization responsibly, safely, impactfully, and effectively. This guide scratches the surface of our experience and understanding of working with organizations of all sizes on their AI use cases and needs.

Start a demo today, and see how we can be a partner in your organization's AI journey.

GitLab Software. Faster.