

# AI-DRIVEN DEVOPS: NAVIGATING THE FUTURE OF SOFTWARE ENGINEERING

## Abstract

This whitepaper explores the DevOps Tools landscape and their AI capabilities at different maturity levels across software development life cycle. It uncovers the transformative potential of these AI capabilities of DevOps tools in realizing higher DevOps ROI (Return on Investment) faster and more efficiently. It also highlights how a strategic and centralized DevOps platform engineering approach can better translate the maturity of AI capabilities in DevOps tools into improved developer experience, increased velocity, and dependable software delivery.

# Insights

- As part of DevOps adoption, 80% of large enterprises who have adopted DevOps have low to high DevOps maturity
- AI adoption in DevOps can act as a catalyst to existing automations and enable higher ROI and faster path to automation
- This whitepaper analyzes the DevOps toolchain across the SDLC, offering insights into their AI maturity levels and highlighting use cases that are primed for leveraging AI-driven augmentation.
- It also examines the future potential use cases getting enabled in DevOps by Agentic AI frameworks

## Introduction

With over 277 million subscribers, a leading video streaming service accounted for 15% of internet traffic in 2023. Back in 2008, it had fewer than 5 million customers, and a missing semicolon in one of their programs caused a major service outage. This incident marked the beginning of their transformation into a technology pioneer, leveraging DevOps, microservices, and cloud computing. This is Netflix which has achieved an impressive operational state where, with a team of under 100 engineers and minimal reliance on a network operations center, it manages hundreds of microservices, thousands of daily production changes, tens of thousands of virtual cloud instances, hundreds of thousands of customer interactions per minute, a global customer base in the millions, and billions of time series metrics. Similarly, Amazon scaled its operations to perform 50 million deployments annually (as early as a decade ago) across development, testing, and production environments. Leading tech giants and innovative companies such as Google, Netflix, Uber, Facebook, Twitter, Spotify, and others are reaching comparable levels of scale and operational speed

How do these companies that handle millions of customers and execute thousands of transactions per second maintain consistently high performance, near-perfect availability, and exceptional customer satisfaction? The core strength lies in their capability to swiftly update applications in real time by leveraging DevOps principles and practices, such as automation, Continuous Integration and Continuous Deployment (CI/CD) pipelines, and cutting-edge DevOps automation tools. All this ensures that changes are deployed efficiently and reliably aligned with the pace of innovation in the ways of working and accelerating industry-wide DevOps adoption

True to agile principles, small and frequent changes and delivery to value are happening with high deployment frequency – instead of infrequent big bang deployments. DevOps ROI is measured by software development speed and stability and likes of these top performing companies typically fall in the elite category, highlighting the efficiency and resilience of top-tier DevOps teams. Refer below table for metrics details – also known as DORA metrics

Common characteristics of elite performers and what it means for business:

- Higher Deployment frequency → better market adaptability, higher customer satisfaction
- Shorter Lead time for changes → faster response to customer needs
- Faster Time to restore service → stronger company reputation & trust
- Lower Change failure rate → better software reliability

	Elite	High	Medium	Low
Deployment frequency	On-demand (multiple deploys per day)	Once per week to once per month	Once per month to once every 6 months	< 1 in 6 months
Lead time for changes	< 1 hr	Between 1 day to 1 week	Between 1 month and 6 months	< 6 months
Time to restore service	< 1 hr	< 1 day	Between 1 day and 1 week	> 6 months
Change failure rate	0% to 15%	16% to 30%	16% to 30%	16% to 30%

80% of enterprises who have adopted DevSecOps have low to high DevSecOps maturity

Figure 1: DevOps Performance Metrics

## Augmentation of DevOps ROI and maturity using AI capabilities

DevOps tools have become the backbone of modern software development, driving automation and efficiency across the SDLC (Software Development Life Cycle). As organizations strive for faster delivery and higher quality, DevOps tools are evolving to meet these demands. The integration of AI features into DevOps tools has been a significant trend for the last 18 to 24 months and it covers entire lifecycle (Refer Figure 2.)




Figure 2: SDLC view

It has opened a world of opportunities to address key challenges and new use cases which were not possible in non-AI world. For example:

Challenge	AI enabled Solution	Business impact
Steep learning curve in adopting new tools	Natural language driven interactions, automated documentation generation and guided workflows	Reduced learning curve Faster and mature adoption of new tools increased productivity
Manual monitoring and troubleshooting of CICD pipelines, deployments, etc.	AI driven monitoring tools – predict failures, automate root cause analysis and suggest resolutions	Faster time to market Better Customer experience
Resource optimizations – scaling up or down infra requirements intelligently	Prediction of workload trends including peak scenarios timings, recommendation or auto-optimize dynamically and in real time	Cost savings Improved performance during peak usage
Limited or no insights into release predictability from successful user adoption, timelines and quality perspectives of new features	Continuous insights driven feature iteration planning based from past trends, failure patterns and sentiment analysis of user base	Improved customer confidence Predictable cost of release

One of the effective strategies and approach we recommend for the organizations is to utilize the inbuilt AI capabilities provided by their DevOps toolchain adopted in the organizations’ ecosystem – that will help augment their DevOps ROI faster and efficiently. This approach can be adopted by the enterprises who are shifting left from low DevOps ROI to Elite or High DevOps ROI. The current maturity of AI in DevOps is low and will be led more by proprietary DevOps tools and large organizations (many of which are our clients) are adopting those capabilities. Additionally, organizations need to make a shift to centralized DevOps tooling platform to better gain value from AI. For experimentation, a few organizations are also leveraging AI features in open-source as proof of concepts in smaller teams and targeted use-cases.

Our recommended AI augmentation approach is based on following key analysis and observations.



## AI capabilities landscape of DevOps tools

Leading DevOps tools and platforms like Atlassian, Microsoft, Harness and Gitlab have already started to release out of the box AI capabilities. These are the market leaders as identified by 2024 Gartner® Magic Quadrant™ for DevOps (refer figure 3. below) and are setting the industry trend in AI enabled tool ecosystem.



Figure 3: Gartner Magic Quadrant for DevOps Platforms





These tools are incorporating AI to optimize continuous integration and continuous delivery (CI/CD) processes, making them hyper efficient and reliable. AI features from the top vendors are getting released at regular intervals with a defined roadmap covering all phases of SDLC enabling richer automation, smarter interactions, value added predictions and augmented assistance


As an example, we surveyed our existing enterprise clients (~80) for Atlassian AI usage for requirements engineering phase and found that 50% of respondents have this feature available, though majority of them are still to adopt it – even though it is bundled at no extra cost for the Enterprises. In some cases (~10%), there is a concern related to security and AI features are not turned on. These inbuilt AI capabilities can enhance productivity by 25 to 30% and standardize and improve the quality of requirements/stories. Some of the impactful use cases include, quickly find issues using natural language rather than creating complex queries, creating existing automation faster using natural language, faster story grooming by recommendations on suggested tasks and sub-tasks, summarize Information/content of a large story or comments on the story, acting as a virtual teammate summarizing

5-page Post-incident Review (PIR) reports into 5-sentence recaps for monthly meetings, etc.

Harness, a leading and popular DevOps and Continuous Delivery platform has defined and implemented a well-defined initial Harness AIDA (AI Development Assistant) roadmap of use cases across SDLC areas. These cases cover automatic creation of Continuous Integration Pipeline, unit test generation, suggested remediations for security vulnerabilities, etc. All of this in a single platform and tool helping to improve developer experience and usability as well.

Open-source tools like popular Continuous Integration tool and orchestrator - Jenkins are enabling AI features via open-source community and in the form of plugins. K8sGPT is another example which uses AI models to reveal problems in Kubernetes cluster and suggest mitigations.

In a short time after its launch, GitHub Copilot reached approximately 1 million+ paid subscribers, 50K+ businesses and is the most widely adopted AI developer tool. 1 in 3 Fortune 500 companies use it. These adoption figures are much higher than many of its open-source counterparts providing AI features.





## AI use cases maturity in DevOps tools across SDLC

We did a holistic analysis of various tools used by our Enterprise clients to understand their AI capabilities that address various DevOps use cases across SDLC and their current and future maturity level. The aim is to provide a benchmark for organization on where they can derive maximum value from AI capabilities of their current DevOps tools stack

- For determining maturity across SDLC, we analysed 30+ industry leading DevOps tools which are widely adopted by our clients and took into considerations various client conversations and requirements around AI driven use cases.
- Low, Medium and High represents the coverage of AI use cases and features in these tools across Day 0, Day 1 and Day 2 activities of SDLC

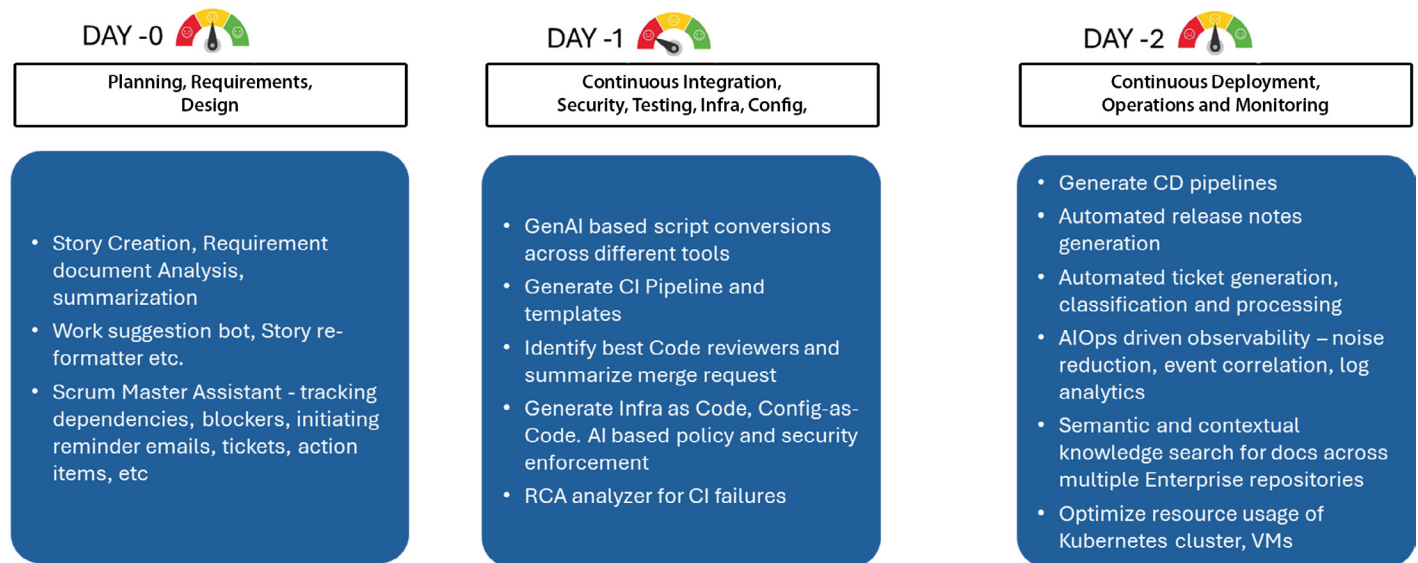


Figure 4: AI Use Cases Maturity Across SDLC



## Key takeaways from analysis on AI maturity are as follows:

Overall maturity is low in DevOps as it is bound to tools and these tools need integration to cover whole of SDLC

Day 0 and Day 2 have higher current and future maturity in AI use cases compared to Day 1 scenarios

DevOps Tools still need to mature in providing Day 1 AI use cases as that is typically the most effort intensive area of SDLC (typically 40% to 50%) and where most interactions happen. That is where Platform Engineering approach can help to bring further Maturity in AI adoption by covering end to end Day 0, 1, 2 AI use cases.

Our observations from Day 0, 1, 2 with respect to AI maturity analysis. Representative clients quoted below include: *Specialty insurance company* | *US based financial group* | *World's leading drug company* | *Leading Cloud based Identity and Access Products and tools company* | *One of world's largest energy company*

### 1) Day 0 - Planning, Requirements and Design:

- Maximum AI implementation in text generation, summarization, and transformation.
- Scrum Master Assistant - handle activities like tracking dependencies, blockers, initiating reminder emails, tickets, action items, etc.
- Auto story grooming, enhancing & standardization of requirements, recommendations on effort estimations. DevOps tools evaluation, migration, setup – standardization, optimizations, modernizations (e.g. Script conversions from one tool to another (e.g. Chef to Ansible)

### 2) Day 1 - Continuous Integration, Testing and Security:

- Focus on code generation, categorization, and suggestions.
- Challenges include tool dependencies, security constraints, and proprietary tool limitations.
- Potential for automating requirement-based templates and customized frameworks via AI

### 3) Day 1, Day 2 - Continuous Deployment and Security

- Least AI penetration in this area - especially in implementing GitOps, Rollback, Validations using AI because of applications constraints like lacking in containerizations, dependencies on custom developments and overall cost of refactoring.
- Optimize resource usage of Kubernetes cluster by anticipating usage based on historical data, trends, patterns and auto-scaling up / down
- AI based policy enforcement – Generate based on best practices and application context and report violations

- IaC code generation – Terraform scripts, Sentinel and Rego policy code scripts
- Future potential is there as more and more legacy apps will get modernized and will adopt latest DevOps tools where AI can be implemented easily

### 4) Day 2 - Continuous monitoring and Site Reliability Engineering

- Evolving solutions for automatic ticket generation, issue aggregation, classification, anomaly detection, and auto-remediation.
- Semantic and contextual knowledge search for docs across multiple Enterprise repositories
- Automated release notes generation
- Future possibilities include near-complete AI driven automation leading to a state of NoOps

### 5) Day 0, 1, 2 – Upcoming trend – Platform Engineering approach (building centralized DevOps Internal Developer Platform or IDP) - Maximum potential of AI driven value adds:

- **DevOps is moving towards centralized Internal Developer Platform** – set of DevOps tools, workflows and technology – **integrated and orchestrated** that enables developer self-service.
- Today, developers manage ~14 different tools on an average for automation and this is a major area impacting developer experience and slowing down the flow of business value. IDPs are emerging as the right-fit solution for this issue. As per **Gartner's hype cycle** in 2022 and top strategic tech trends in 2023, 2024, 80% of software engineering organizations will establish platform teams by 2026

- Integration of generative AI to IDPs has high degree of value add due to the rich data and insights from all DevOps tools available at one place – in addition to enhance self-service capabilities, it can also address customized use cases, along with better traceability, centralized tooling, security controls, compliance and governance, under single umbrella at enterprise level. Upcoming Agentic AI can disrupt and address multiple use cases autonomously via centralized IDP approach. Refer
- Figure 5 below for our view of an Intelligent Platform Engineering approach to enable wider coverage and penetration of AI value add across Day 0, Day 1 and Day 2 aspects of DevOps. We envision a platform centric approach which centralized and scales specialized knowledge and actions across the entirety of development and operations lifecycle by plug and play Agentic AI agents
- Benefits:** Dev productivity | Costs | Security & Compliance | Time to Business Value | Traceability | Near-autonomous operations

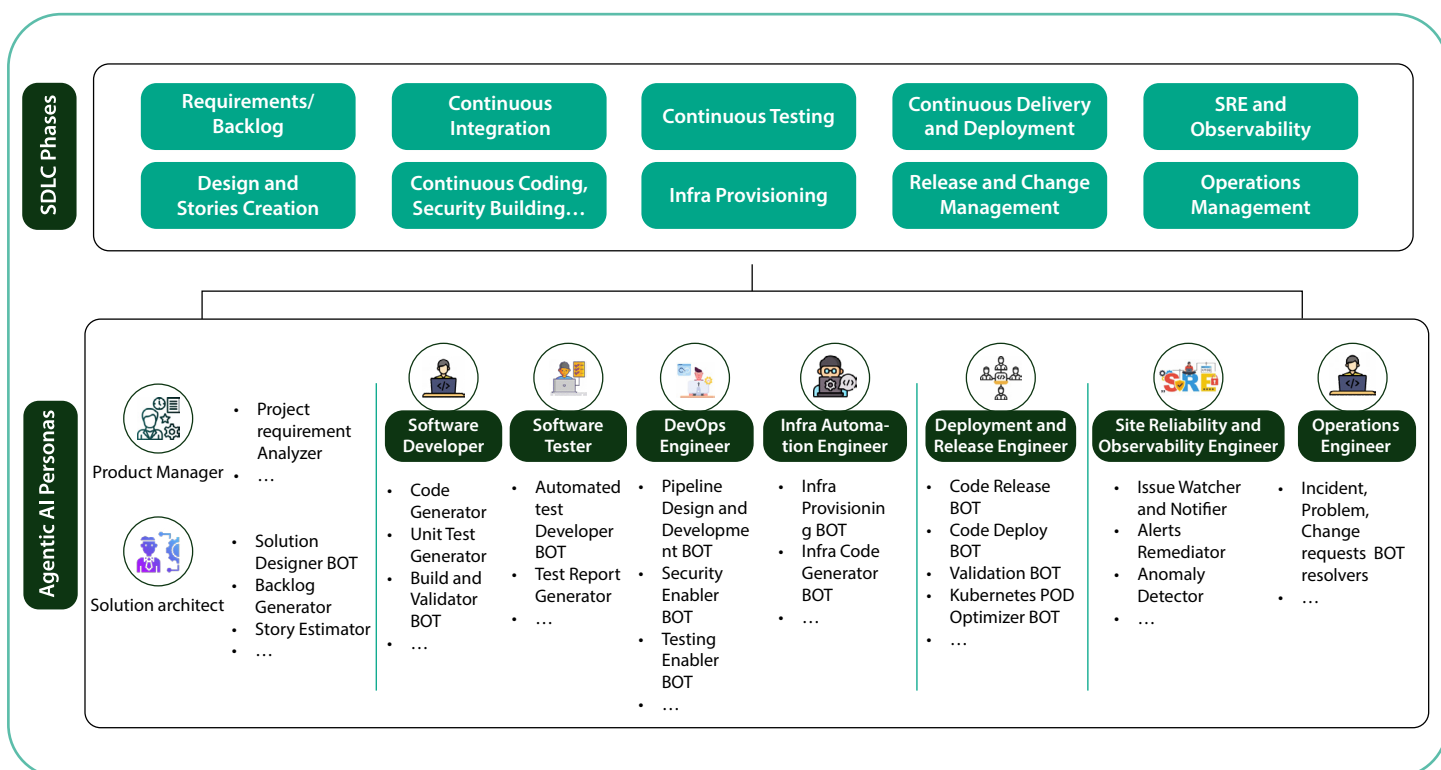


Figure.5. A representative depiction Agentic AI Personas across SDLC

A sample interaction of agents with Platform Engineering approach is provided below (Figure 6) which considers an example of an Infra Automation Engineer persona agents capable of achieving Infrastructure provisioning on demand

- The main agent below can interact with users and get the requirements, Let's say, users have given input to provision Linux virtual machine provisioning.
- As part of formalization agent can ask size, type etc. to user, ensures to get all inputs and format it
- Send the inputs to infrastructure provisioning agents and wake it up to start provisioning without disturbing other agents. Infra agent can further trigger subagents to generate code, enable security, plan and deploy code to provision infrastructure respectively
- Further interact with subagents capable of understanding infrastructure dependencies and update the code as per user requirements.
- Post provisioning records the provisioning information or anomalies identified in knowledge base subsequently enabling learning for next virtual machine provisioning.



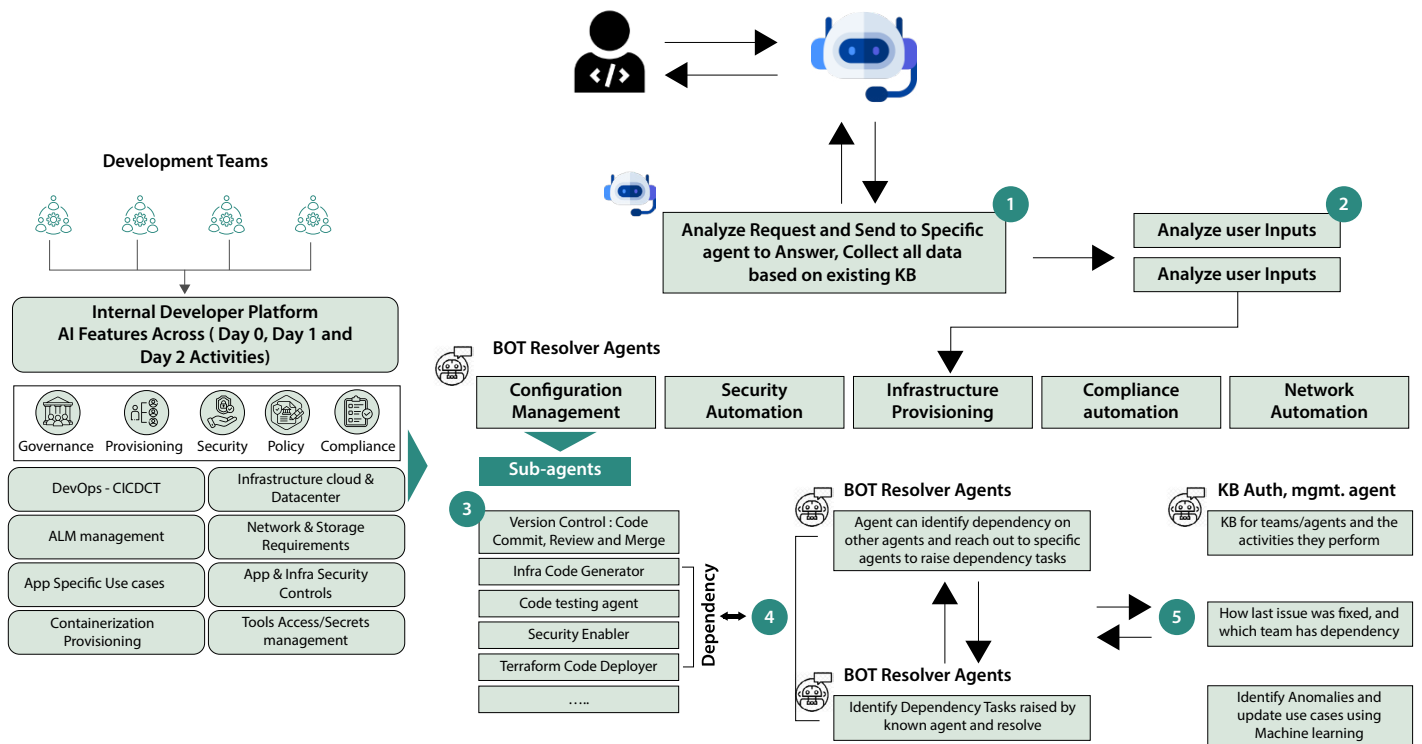
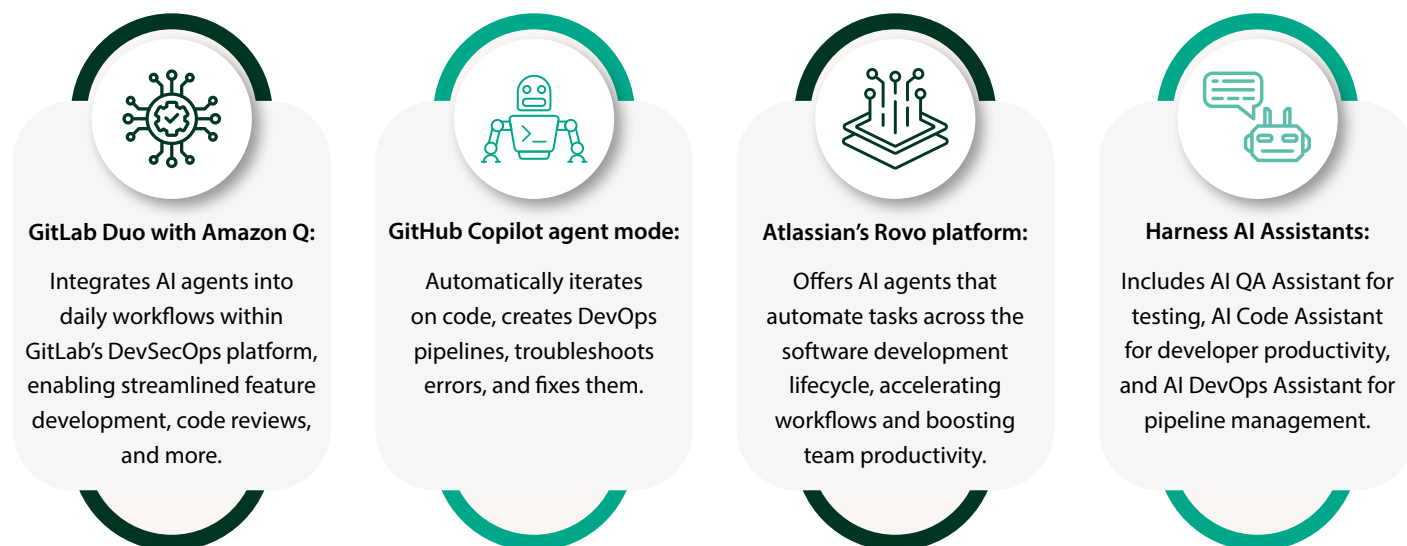


Figure 6: A representative depiction of Agentic AI in action in an Internal Developer Platform



Leading platform vendors are advancing agentic AI to enhance workflows and productivity which give a glimpse of near future possibilities of AI enabled DevSecOps. Key examples include:



## Factors to be considered for adoption of AI enabled DevOps tools

Based on our experience with multiple clients, we are observing some initial barriers and challenges in adoption of AI via DevOps tool. Following are some of the key factors to be considered while adopting AI features of DevOps tools



## Conclusion

DevOps drives SDLC automation, achieving exceptional ROI as demonstrated by industry leaders. Organizations aiming to elevate their performance from high to elite can leverage AI capabilities within DevOps tools, accelerating transformation.

Adopting an AI-augmented approach and utilizing intelligent, centralized DevOps platforms with specialized agents can enhance workflows to achieve autonomous or NoOps systems. Such AI driven platforms also unify enterprise knowledge by consolidating data across DevOps tools for planning, development, testing, infrastructure, security, compliance, and operations.

To maximize ROI, organizations should select tools tailored to their specific challenges and objectives, rather than retrofitting AI tools into existing systems. In many cases, developing custom AI agents aligned with enterprise needs and principles can complement the native AI features of their tools landscape effectively.

## Key Definitions

- **Day 0, Day 1, Day 2:** The Software Development Life Cycle (SDLC) functionality can be divided into Day 0, Day 1, Day 2. The concept of days aligns with agile principles of software development like iterative, early and continuous feedback, continuous improvements and collaboration.
- **Day 0:** Planning and design activities where many collaborations related DevOps tools are utilized
- **Day 1:** Activities where tools are used and integrated to enhance developer productivity doing continuous integration, code testing and security, Infra, policy, config as code and continuous deployment
- **Day 2:** Activities which covers release management, change management, logging, monitoring and application and DevOps infrastructure operations management
- **DORA Metrics:** DevOps Research and Assessment (DORA) offers a standardized set of metrics for assessing DevOps process performance and maturity. These metrics highlight the speed of response to changes, average code deployment times, iteration frequency, and insights into failures.
- **Platform Engineering:** Platform engineering focuses on creating and developing Internal Developer Platforms (IDPs) that empower developers with self-service capabilities while ensuring standardization is embedded into the design process.
- **IDP:** Internal Developer Platforms. Self-serve platforms to provide a curated set of tools, reusable capabilities for different DevOps tools and processes packaged for easy consumption by end user personas with application of all guardrails & governance within an enterprise
- **laC:** Infrastructure-as-Code. Managing and provisioning of infrastructure through code





## References

In developing this whitepaper, valuable information and insights were gathered from a variety of trusted sources, such as research papers, articles, and other resources. Key references that shaped the content of this whitepaper include:

1. Demystifying Microservices with Node.js: A Journey from Netflix's Semicolon Crisis to the Future of Software <https://dev.to/davydocsurg/demystifying-microservices-with-nodejs-a-journey-from-netflixs-semicolon-crisis-to-the-future-of-software-103e>, <https://www.sandvine.com/inthenews/netflix-is-responsible-for-15-of-global-internet-traffic-consumption#:~:text=According%20to%20The%20Global%20Internet,internet%20traffic%20usage%20in%202022>
2. Google: Elite DevOps performer <https://cloud.google.com/blog/products/devops-sre/using-the-four-keys-to-measure-your-devops-performance>
3. Gartner. Magic Quadrant for DevOps Platforms. Gartner, Inc., 2025. [Online]. Available: <https://www.gartner.com/en/documents/5728483>
4. [GitHub Co-pilot: Your AI Programmer](#)
5. Gartner: Hype Cycle <https://www.gartner.com/en/articles/what-s-new-in-the-2022-gartner-hype-cycle-for-emerging-technologies>
6. Gartner: Platform Engineering That Empowers Users and Reduces Risk <https://www.gartner.com/en/infrastructure-and-it-operations-leaders/topics/platform-engineering#:~:text=Platform%20Engineering%20That%20Empowers%20Users,improving%20the%20experience%20of%20users.>
7. Platform Tooling | Platform Engineering <https://platformengineering.org/platform-tooling> <https://platformengineering.org/blog/ai-is-changing-the-future-of-platform-engineering-are-you-ready>
8. Boston Consulting Group: The Art of Scaling GenAI in Software <https://www.bcg.com/publications/2024/the-art-of-scaling-genai-in-software>

## About the Authors



**Amit Sadana**

Industry Principal, Infosys  
Engineering Services



**Satya Parkash Pandey**

Principal Consultant, Infosys  
Engineering Services



Infosys Topaz is an AI-first set of services, solutions and platforms using generative AI technologies. It amplifies the potential of humans, enterprises, and communities to create value. With 12,000+ AI assets, 150+ pre-trained AI models, 10+ AI platforms steered by AI-first specialists and data strategists, and a 'responsible by design' approach, Infosys Topaz helps enterprises accelerate growth, unlock efficiencies at scale and build connected ecosystems. Connect with us at [infosystopaz@infosys.com](mailto:infosystopaz@infosys.com).

For more information, contact [askus@infosys.com](mailto:askus@infosys.com)

**Infosys**<sup>®</sup>  
Navigate your next

© 2025 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.