theCUBE research

chronosphere

Google Cloud

Scaling Cloud-Native Applications with CI/CD, AI, and Secure Developer Tooling

Day 0 Build, Test Survey Research Report

Paul Nashawaty
Principal Analyst | theCUBE Research

September 2025

Executive Overview

In theCUBE Research's Application Development and Modernization Day 0 Build, Test research survey, we capture how organizations prepare applications before code reaches production. The report analyzes cloud-native architecture readiness, CI/CD and automation maturity, API management, developer tooling, GitOps adoption, and AI integration based on responses from 375 professionals in the fields of IT, development, operations, and platform engineering.

Findings show that cloud-native adoption has become mainstream, with containers and Kubernetes widely used. Yet complexity, skills shortages, and uneven automation still hinder scaling and efficiency. Observability and AI-driven insights are emerging as critical enablers that lower MTTR and boost developer productivity. This report highlights the main issues, new best practices, and areas for leadership action.

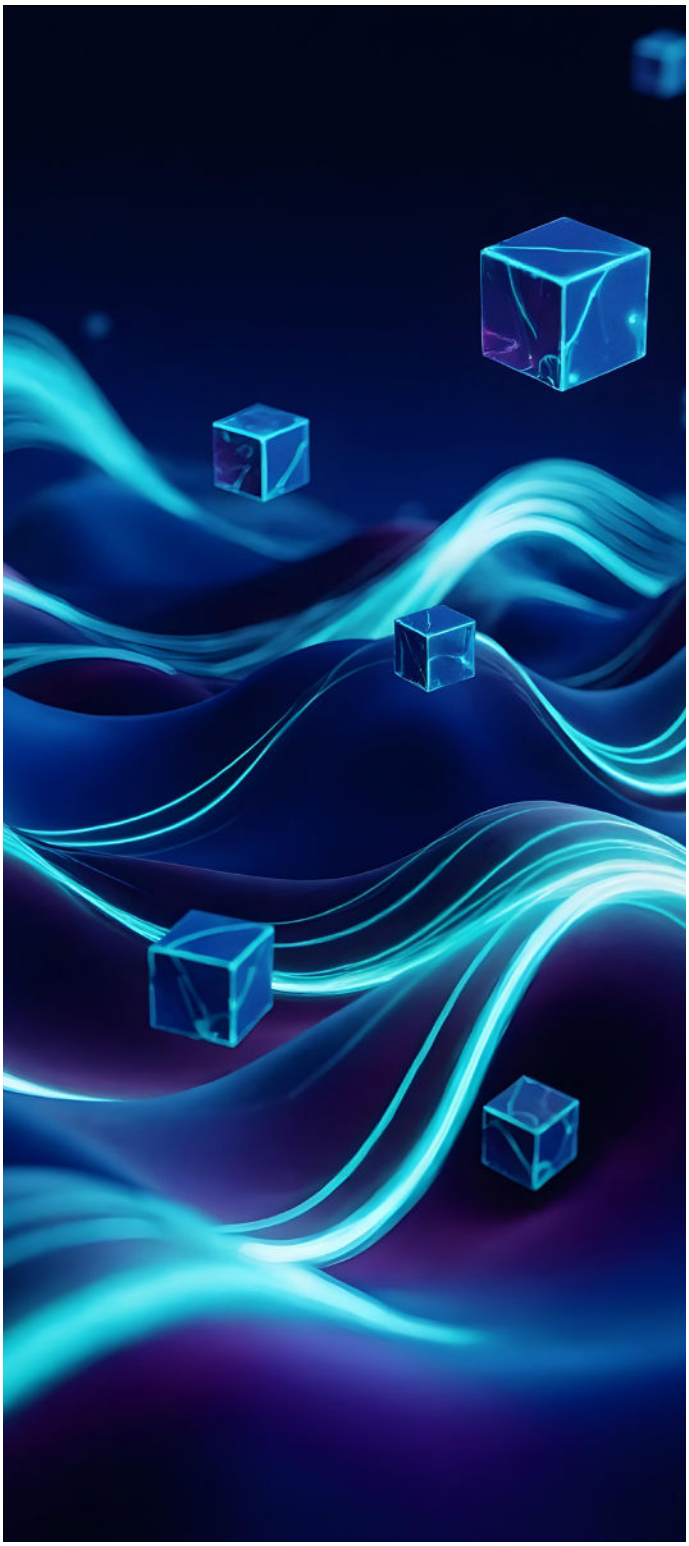

Cloud-Native Readiness

Cloud-native workloads have become commonplace. Nearly half of organizations report that 51–75% of their applications run in containers, representing a sharp rise from just a few years ago.

Chart 1: Containerization Adoption Levels

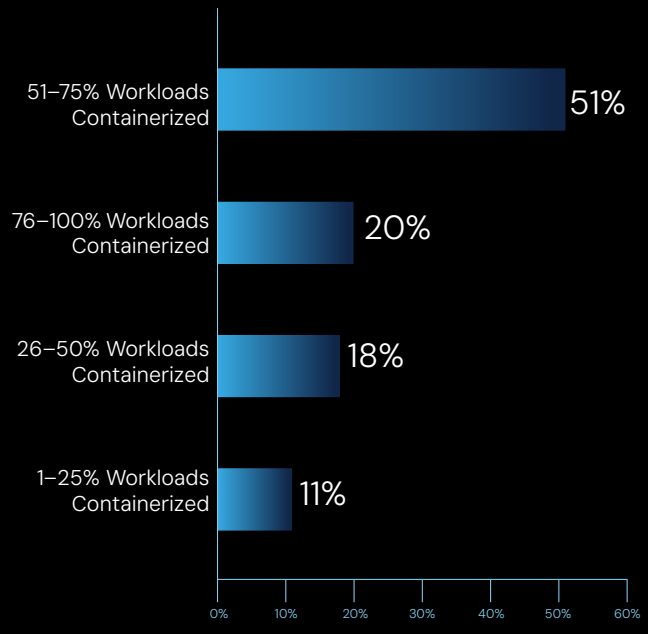

Managed Kubernetes services, such as Anthos, AKS, and EKS, dominate deployment choices, signaling a preference for platforms that integrate with ecosystems and reduce operational burdens. Future advantage will depend on optimizing orchestration, runtime security, and policy enforcement across clouds. Teams must use cross-cloud policy control and interoperability to strike a balance between flexibility and governance.

Developer Tools and API Management

Nearly one-third of respondents still handle versions by hand, despite 89% of them maintaining a centralized API repository. This leads to slower release cycles, governance problems, and version drift risks.

Chart 2: API Management Practices

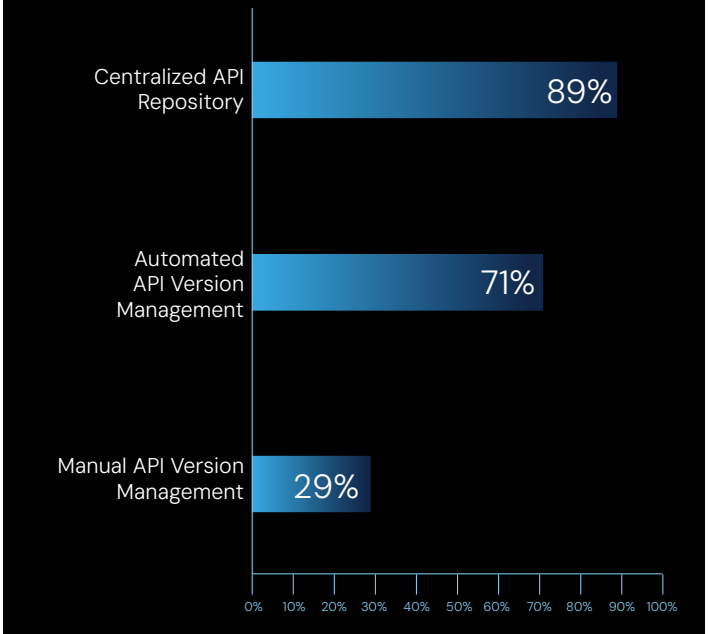

Automation in API lifecycle management is critical to scalability and reliability. Opportunities for unified solutions that integrate governance, observability, and version control are highlighted by poorly integrated toolchains, which slow delivery and increase deployment errors.

Automation & CI/CD Adoption

CI/CD pipelines are widely used, but their maturity varies. Only 27% of organizations validate daily, compared to 51% of organizations that test weekly. This lower cadence increases the chance of undetected issues that delay releases or surface post-deployment.

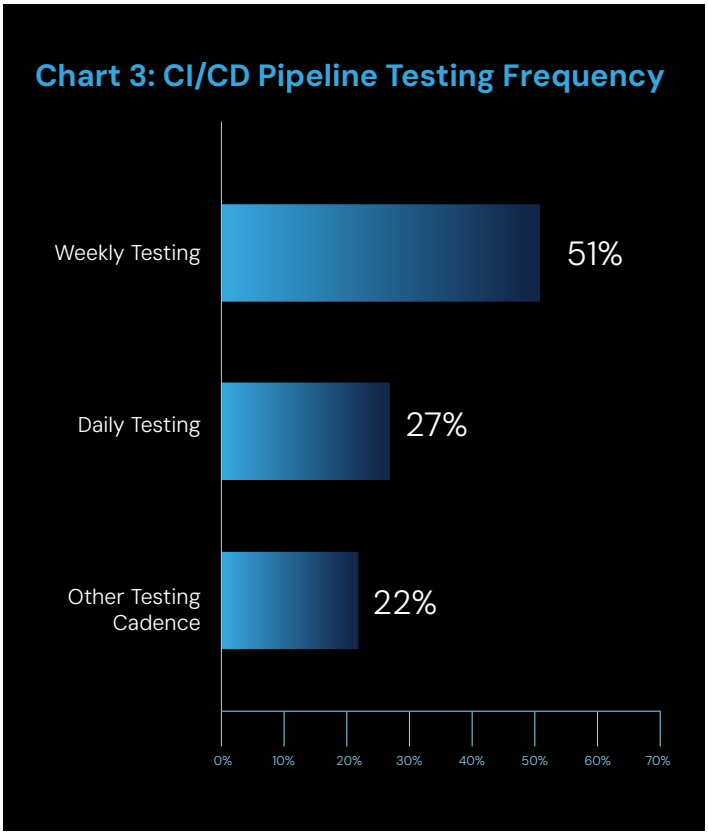

Code scanning adoption is stronger, with 84% already using scanning tools and another 9% planning to adopt them.

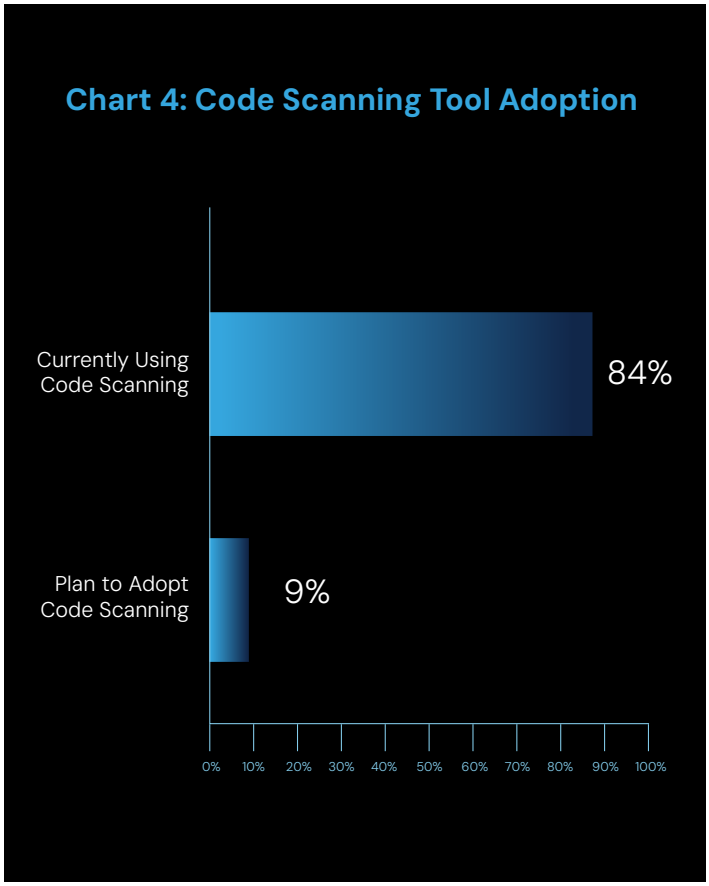

According to the data, increasing test frequency and extending automation coverage are key to efficiency gains. Reliability is enhanced through daily or per-commit testing, and remediation is expedited by integrating scanning tools with developer education. The lack of automated rollback remains a key risk, leading to prolonged downtime when releases fail.

Observability, Security & AI Integration

Application environments produce more telemetry than ever. To manage this, 85% of organizations now utilize AI-powered tools for real-time issue detection, while the remaining 15% plan to adopt them.

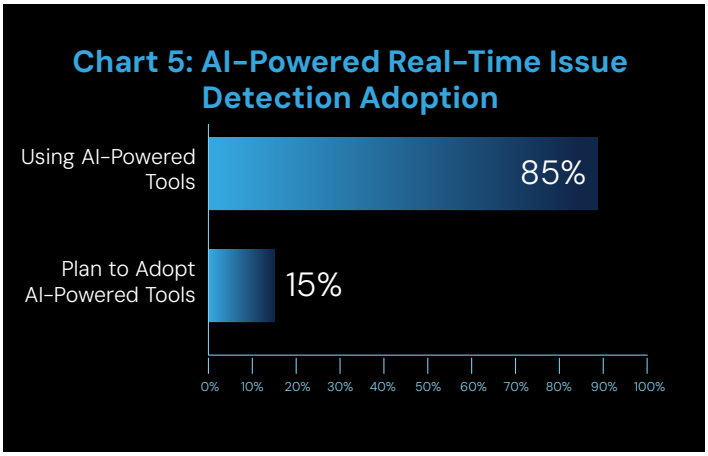

Additionally, security practices are also evolving: 51% identify container image vulnerabilities weekly, and 27% identify them daily, an improvement from older patterns where vulnerabilities often lingered.

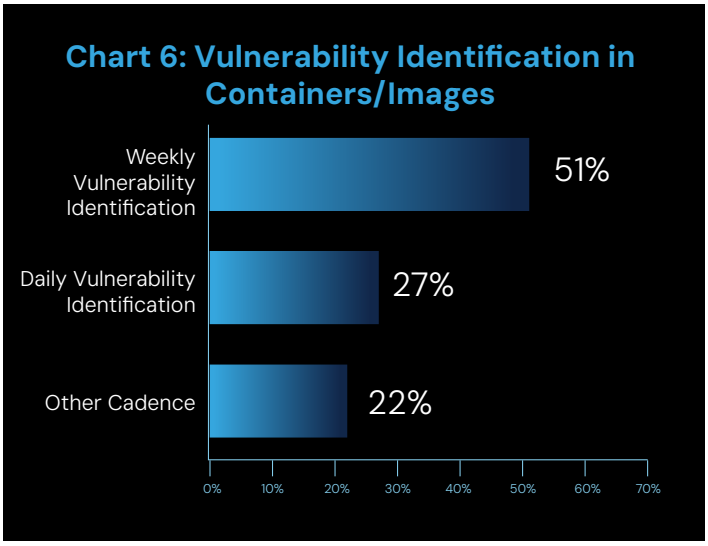

Budget priorities support this trend. With a quarter of most enterprise IT budgets allocated to AI initiatives, half of organizations spend between 26% and 50% of their IT budget on application development. Furthermore, 70% of respondents say AI projects will be their top priority in the upcoming year, followed by cloud infrastructure (60%) and security and compliance (65%). DevSecOps tooling, API management, and developer tools remain lucrative investment areas.

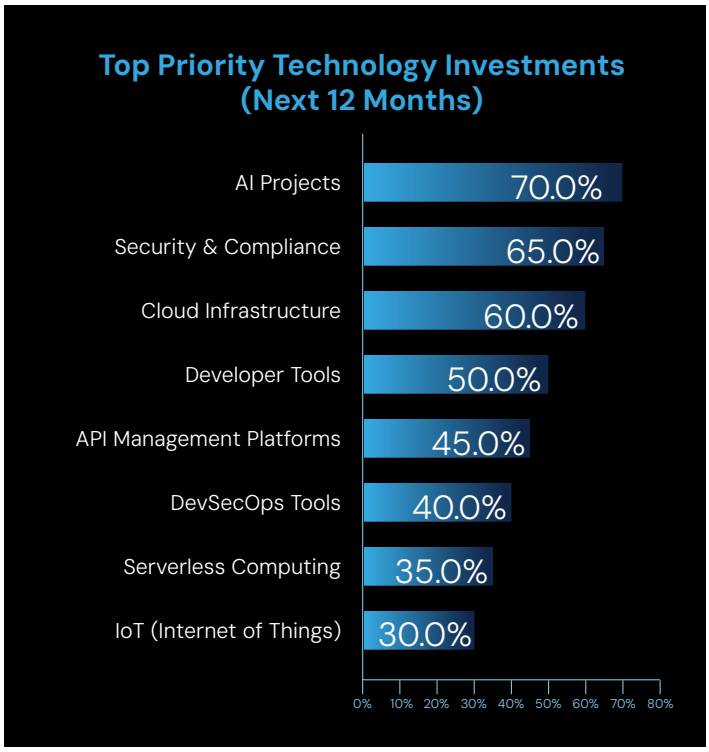

These findings demonstrate the increasing importance of AI-driven observability for ensuring performance and uptime. While budget trends show an industry doubling down on AI, security, and infrastructure modernization, regular vulnerability scanning is now standard practice.

Taming Cloud-Native Complexity: Why Smarter Observability Is the Key to Reliability and Efficiency

For most of the last decade, observability has been about scale—more telemetry, more dashboards, more data. As systems became cloud-native and distributed across containers and clusters, the instinct was to collect everything and figure it out later. That worked for a while. But as recent industry research shows, many teams are now hitting a wall: they have more information than insight, more metrics than meaning.

Today, nearly half of organizations run most of their workloads in containers, often across managed Kubernetes environments. The benefits are undeniable—speed, scalability, and resilience—but so are the challenges. Complexity has become the new constraint. Every deployment, every API version, every new integration adds another layer to monitor and manage. The question is no longer, "how much data can we capture?" But, "how confidently can we understand what it tells us?"

From Collection to Comprehension

According to recent findings, 89% of organizations keep a centralized API repository, yet almost a third still manage versions manually. Testing cycles tell a similar story: only a quarter of teams test code daily, while over half do so weekly. These gaps create blind spots between innovation and reliability, where undetected issues slip into production and slow down recovery.

This is where the story of observability is evolving—from a discipline of collection to one of comprehension.

The shift isn't just technological; it's cultural. Teams are recognizing that effective observability isn't about watching everything, it's about knowing what matters. The goal is no longer dashboards full of noise but guided visibility—where telemetry, topology, and time combine to form understanding.

The Next Phase: Intelligence with Context

AI is already part of this transition. The research shows 85% of organizations now use AI for real-time issue detection. But AI alone isn't a silver bullet. Without context, it risks becoming just another source of noise. The future of observability isn't AI that replaces engineers; it's AI that reasons with them—tools that surface the most relevant signals, generate informed hypotheses, and guide human judgment instead of overwhelming it.

Trust is the currency of this shift. To earn it, observability must be transparent, showing its reasoning and evidence, not just its results. Confidence grows when systems explain why they think something's wrong, not just that something is. In time, these guided insights evolve into agentic workflows—where human oversight meets machine efficiency, and autonomy expands only as understanding deepens.

Observability as a Strategic Discipline

The same research found top investment priorities in AI, security, and cloud infrastructure modernization. Observability sits at the intersection of all three. When it's done well, it becomes a foundation for reliability, performance, and cost efficiency. It's how organizations turn complexity into control and maintain confidence as their environments scale.

The observability journey is no longer about collecting every metric—it's about connecting them. In this new era, the real signal is trust: trust in data, in context, and in the systems that help us see the truth within both.

Conclusion

The 2025 Day 0 research survey reveals insights into a maturing cloud-native ecosystem, where containerization, Kubernetes, CI/CD pipelines, and AI-powered observability are now foundational to enterprise application delivery. While adoption has accelerated, persistent gaps in automation, governance, and security practices continue to challenge scale and reliability.

The findings further underscore that organizations must strike a balance between speed and control, integrating modern developer tooling, DevSecOps practices, and AI-driven insights to reduce risk and enhance productivity. As cloud-native becomes the new standard, leadership focus should shift toward unifying toolchains, embedding automation throughout the pipeline, and aligning budgets with emerging priorities such as AI and security. To dive deeper into the full dataset, benchmarks, and practitioner insights behind this research, please reach out to theCUBE Research for direct access and tailored advisory services.

Disclaimer

Contact

Silicon Valley
989 Commercial Street
Palo Alto, CA 94303

Boston Metro
95 Mount Royal Avenue
Marlborough, MA 01752

David Butler
david.butler@siliconangle.com
774-463-3400