CLOUD NATIVE
COMPUTING FOUNDATION

THE LINUX FOUNDATION | Research

CNCF Annual Cloud Native Survey

The infrastructure of AI's future

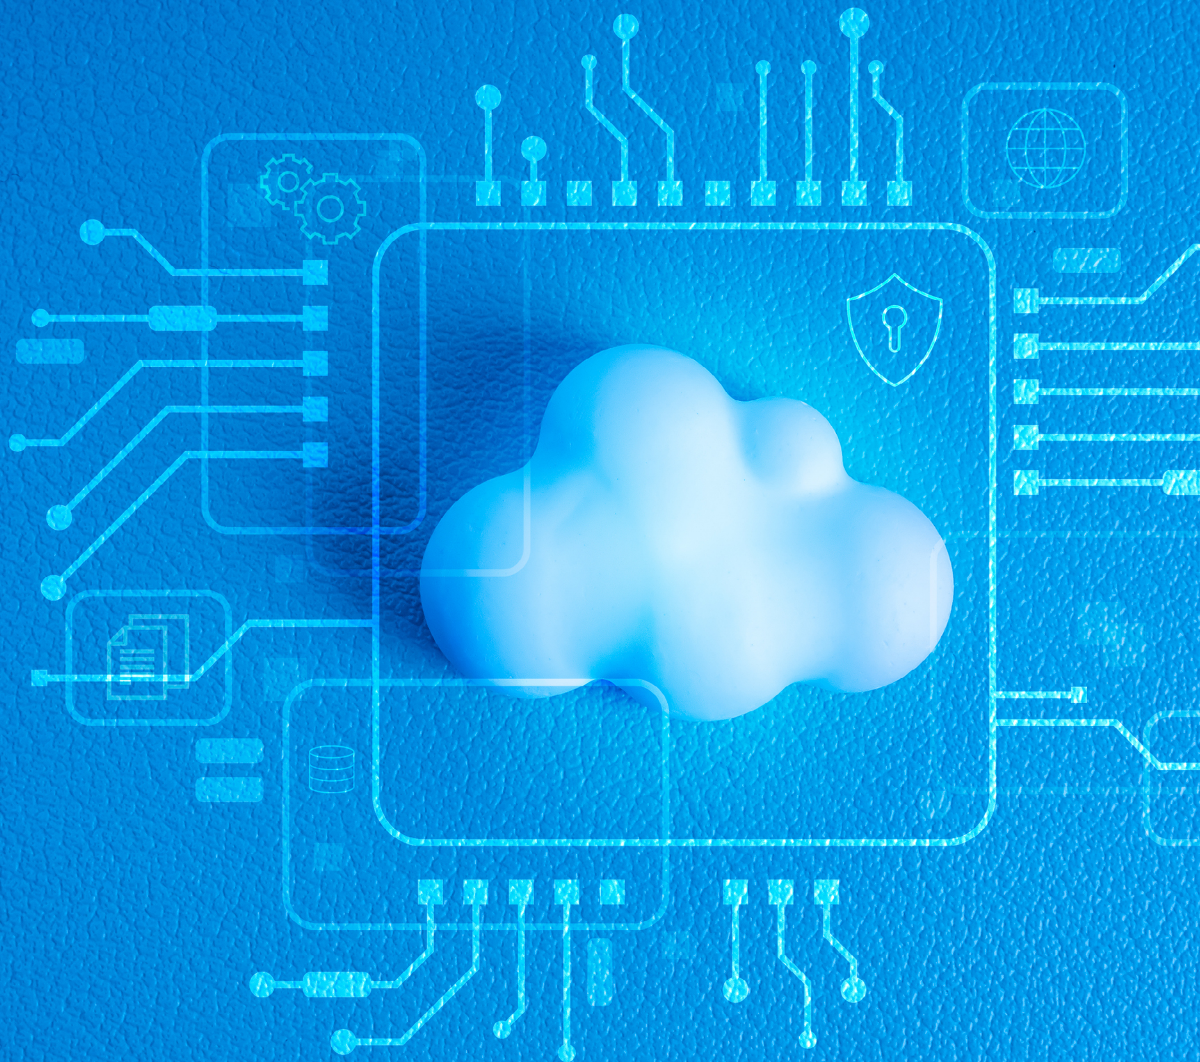

Adrienn Lawson, *The Linux Foundation*

Jeffrey Sica, *Cloud Native Computing Foundation (CNCF)*

Foreword by Jonathan Bryce, Executive Director, Cloud and Infrastructure, *The Linux Foundation*

January 2026

CNCF Annual Cloud Native Survey: The infrastructure of AI's future

CNCF ANNUAL CLOUD NATIVE SURVEY

66% of organizations use Kubernetes to host generative AI workloads.

CNCF ANNUAL CLOUD NATIVE SURVEY

82% of container users deploy Kubernetes in production, up from 66% in 2023.

CNCF ANNUAL CLOUD NATIVE SURVEY

Cloud native adoption reaches 98% across surveyed organizations.

CNCF ANNUAL CLOUD NATIVE SURVEY

Cultural changes with the development team ranks as the #1 challenge for deploying containers (47%), surpassing technical concerns.

CNCF ANNUAL CLOUD NATIVE SURVEY

Container usage in production applications increased from 41% (2023) to 56% (2025).

CNCF ANNUAL CLOUD NATIVE SURVEY

47% of organizations deploy AI models occasionally, with only 7% deploying daily.

CNCF ANNUAL CLOUD NATIVE SURVEY

52% of organizations don't build or train their own AI models - they're consumers.

CNCF ANNUAL CLOUD NATIVE SURVEY

GitOps adoption jumps from 0% among cloud native explorers to 58% among cloud native innovators.

CNCF ANNUAL CLOUD NATIVE SURVEY

CI/CD adoption reaches 91% among mature organizations, making it the most fundamental practice.

CNCF ANNUAL CLOUD NATIVE SURVEY

74% of cloud native innovators check in code multiple times daily vs. 35% of explorers.

CNCF ANNUAL CLOUD NATIVE SURVEY

59% of organizations use cloud native for much or nearly all development (up from 54% in 2023).

CNCF ANNUAL CLOUD NATIVE SURVEY

Infrastructure sustainability emerges as a critical concern with machine-driven automated usage.

Contents

Foreword

This year, we celebrate the 10-year anniversary of the Cloud Native Computing Foundation. A decade ago, the concepts we now consider industry standards were a revolutionary vision. Today, they are the undisputed foundation of modern software systems.

I was honored to step into the role of Executive Director midway through this banner year. This report's findings confirm that cloud native has truly arrived. 98% of organizations have adopted cloud native techniques, and Kubernetes usage in production has soared to 82% among container users.

This maturity has revealed our community's next great opportunities. The primary challenge in cloud native is no longer technical complexity; 47% of organizations now cite "cultural changes with the development team" as their top obstacle. This is the "why" behind Platform Engineering. Our community is leading this charge, building the standard platforms that tame complexity and unlock developer velocity. We can see this in GitOps adoption among other practices. Across industries, teams are productizing the path from code to production with paved roads, sensible defaults, and clear guardrails. The payoff is less undifferentiated repetition for developers and faster, safer delivery for the business.

This stable, platform-centric foundation is also enabling the next major shift impacting the entire industry: Artificial Intelligence.

This report is titled "The infrastructure of AI's future," and its data reveals the emerging substrate: 66% of organizations are already using Kubernetes to host their generative AI workloads. But the real story isn't the one in the headlines. It's not about training LLMs. Most enterprises do not build or train their own models—they are consumers. The real challenge is deployment, and the gap between ambition and reality is stark.

If we cut through the hype of chatbots and agents, we can clearly see that we will need to greatly decrease the difficulty of serving AI workloads while massively increasing the amount of inference capacity available across the industry. I believe this is the next great cloud native workload. Our community has the experience and expertise to run, scale, and observe models–all while openly sharing the technology and practices that enable global innovation.

This new wave of AI also brings a new responsibility. As the report warns, "machine-driven automated usage" puts a strain on the open source systems we all depend on. Our next decade must be defined not just by innovation, but by continuing to evolve and support our incredible contributors.

Thank you to everyone who contributed code, documentation, testing, stories, and support this year. If you're new to CNCF, welcome; you'll find a community that turns hard problems into shared solutions. If you've been with us since the early days, thank you for building the foundation we continue to strengthen—together.

Onward to the next decade of cloud native!

JONATHAN BRYCE,
Executive Director, Cloud and Infrastructure,
The Linux Foundation

Executive summary

The 2025 CNCF Annual Report reveals a cloud native ecosystem that has reached a critical inflection point. What began as an experimental architectural approach has solidified into an enterprise infrastructure standard, with 98% of organizations now employing cloud native techniques. The story, however, is no longer about technology adoption. It's about maturity, sustainability, and the quiet revolution happening beneath the AI hype cycle.

Three key themes define 2025:

First, Kubernetes has evolved from container orchestrator to AI infrastructure platform, with 66% of organizations running generative AI workloads on it. The primary obstacles have shifted from technical complexity to organizational transformation, with cultural resistance now the leading challenge at 47%. The sustainability of open source infrastructure has emerged as an existential concern, as machine-driven automated usage strains essential systems that help build, test, deploy, and distribute software.

Second, the data demonstrates that cloud native maturity follows a predictable progression model. Organizations advance through four distinct stages: explorers, adopters, practitioners, and innovators. Each characterized by specific technology adoption patterns and development velocities. GitOps serves as a north star metric: not one of the explorers have implemented it, while 58% of innovators run GitOps-compliant deployments.

Third, the report reveals a profound gap between AI ambition and infrastructure reality. While headlines concentrate on model breakthroughs, 47% of organizations deploy AI models only occasionally, and 52% do not train models at all. The real competitive advantage lies not in algorithms but rather in the unglamorous infrastructure capabilities: robust CI/CD pipelines, and resource optimization.

Introduction

Cloud native technologies have become fundamental to modern enterprise infrastructure. As organizations deploy AI workloads, modernize applications, and pursue operational agility, understanding actual adoption patterns and implementation practices is critical for informed decision-making.

This report presents findings from the 2025 CNCF Annual Survey, which examines organizational perspectives on containers, Kubernetes, and cloud native adoption trends. The survey collected data from 628 respondents across multiple industries, company sizes, and geographies during September 2025. As an annual study, this survey enables year-over-year tracking of technology adoption and maturity, revealing not just current practices but also trajectory and evolution of the cloud native ecosystem.

The Cloud Native Computing Foundation (CNCF) ecosystem now comprises 234 projects at various maturity levels, supported by over 270,000 contributors from all over the globe. This report analyzes adoption patterns, deployment practices, and emerging trends across the cloud native landscape, providing benchmark data for organizations to assess their own cloud native maturity and strategy.

Kubernetes as the AI platform

While headlines tout AI breakthroughs, a quieter revolution is happening in the infrastructure layer. CNCF survey data reveals that 66% of organizations are betting on Kubernetes to run their generative AI workloads. However, success depends on mastering the unglamorous challenges of resource management and deployment pipelines.

The rise of Kubernetes as the de facto AI platform represents a fundamental shift in how organizations approach machine learning operations. Traditional ML infrastructure often relied on specialized, monolithic platforms that created silos between data science teams and production engineering. Kubernetes bridges this gap by providing a unified orchestration layer that handles both traditional application workloads and compute-intensive AI tasks. Projects such as Kubeflow provide end-to-end ML workflows, while KServe handles model serving at scale. The introduction of GPU scheduling capabilities, node affinity rules, and sophisticated resource quota management allows organizations to efficiently share expensive hardware resources across multiple teams and workloads.

Kubernetes AI adoption wave

As shown in Figure 1, Kubernetes has become the de facto orchestration layer for production AI, but the split between full adoption (23%) and partial adoption (43%) reveals organizations are taking a measured, infrastructure-first approach.

The 23% with full Kubernetes adoption for inference workloads represent organizations that have achieved true MLOps maturity. These teams have likely implemented GitOps workflows for model deployment, established robust monitoring through Prometheus and Grafana for model performance metrics, and integrated AI workloads into their existing CI/CD pipelines.

The 43% partial adoption group typically starts with Kubernetes for specific use cases, often batch inference jobs or development and staging environments, while maintaining legacy systems for production serving. The 18% planning to adopt Kubernetes could be facing obstacles including existing investments in proprietary ML platforms, concerns about operational complexity, or the need to retrain their teams.

Moving AI workloads to Kubernetes isn't simply about containerization. Organizations must address unique requirements such as managing large model artifacts through container registries or object storage. They also need to ensure that models land on appropriately resourced nodes with GPU affinity, architecting different patterns for training pipelines versus low-latency serving, and implementing canary deployments and rollback strategies specifically designed for ML models.

FIGURE 1

KUBERNETES USAGE FOR HOSTING AI WORKLOADS

Do you use Kubernetes to host your generative AI workloads? (select one)

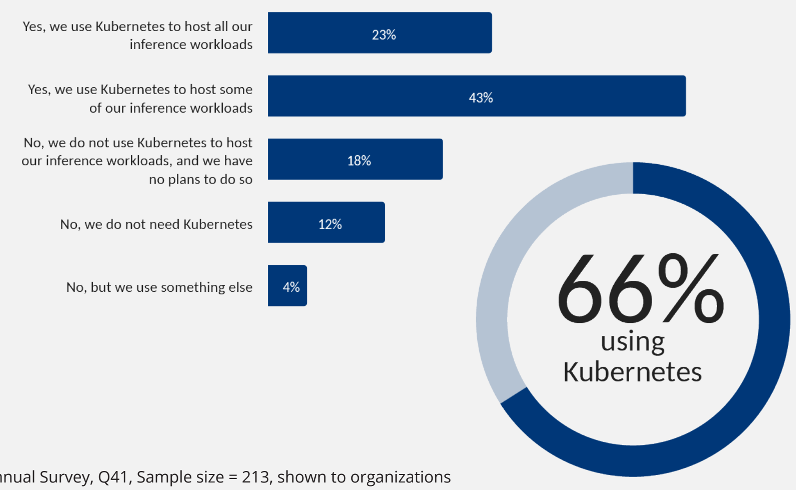

2025 CNCF Annual Survey, Q41, Sample size = 213, shown to organizations who host AI workloads based on Q39 and Q40, DKNS excluded

Most organizations are consumers of AI

The contrast between inference hosting and training tells a crucial story in Figure 2. Most organizations are consumers, not producers, of AI models. 52% of surveyed organizations do not build or train AI models and those who do likely do not build from scratch but rather fine-tune based on their own data. This has implications for infrastructure needs as inference workloads require different scalability and cost optimization strategies.

Organizations consuming pre-trained models face a distinct set of infrastructure challenges. The primary focus shifts to inference optimization through techniques like model quantization, ONNX runtime optimization, and batching strategies. While training requires expensive GPUs for hours or days, inference runs continuously, making cost management critical. Organizations implement sophisticated autoscaling policies, potentially using CPU-based inference for less demanding workloads and reserving GPU resources for latency-sensitive applications.

The 37% using managed APIs represent organizations prioritizing speed-to-market over infrastructure control. However, even these teams benefit from Kubernetes-based orchestration layers that can implement retry logic and fallback strategies across multiple providers, cache common responses to reduce API costs, abstract provider-specific APIs behind unified interfaces, and monitor usage and costs across different services. The 25% self-hosting models have made the economic calculation that ownership costs justify the infrastructure investment, which typically makes sense when request volumes exceed one million inferences monthly, data privacy regulations prohibit cloud API usage, or latency requirements demand local deployment. On an even more local level. Edge deployment (13%) is an emerging pattern requiring specialized orchestration.

FIGURE 2

AI HOSTING

Does your organization host generative AI models for **inference**? If so, where? (select all that apply)

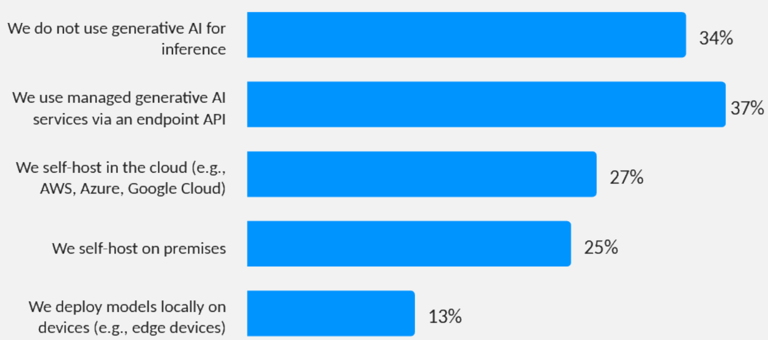

2025 CNCF Annual Survey, Q39, Sample size = 497, DKNS excluded.

Does your organization host generative AI models for **building and training**? If so, where? (select all that apply)

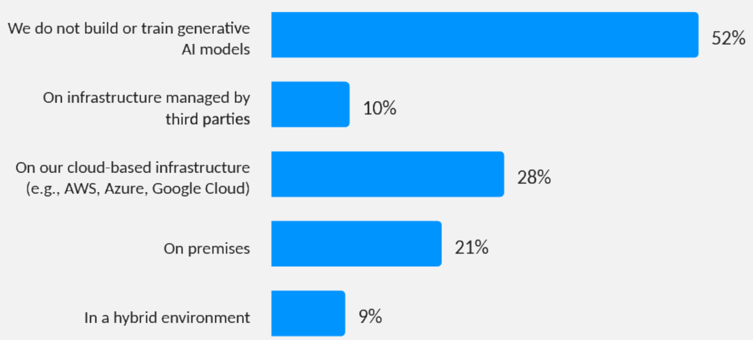

2025 CNCF Annual Survey, Q40, Sample size = 482, DKNS excluded

Deployment maturity

The deployment frequency data shows a reality check (Figure 3). 47% deploy occasionally, only 7% daily. The AI revolution is methodical. It is a production-grade deployment requiring robust CI/CD, monitoring, and governance infrastructure.

Unlike traditional code where unit tests and integration tests provide confidence, AI models require statistical validation through testing model performance on holdout datasets among other complex tests. These validation gates slow deployment velocity but are essential for production reliability.

The small minority achieving daily AI deployments, only 7%, have likely implemented automated retraining pipelines that continuously incorporate new data. These organizations treat models as living systems requiring constant updating rather than static artifacts. The remaining 93% of organizations are far from this state.

Diversity in AI workloads

Among those who run AI/ML workloads on Kubernetes, there's significant diversity in use cases beyond the hype (Figure 4). Real AI/ML adoption is about practical infrastructure challenges, not just buzzwords.

FIGURE 3

FREQUENCY OF AI MODEL DEPLOYMENT

How frequently does your organization deploy generative AI models into production? (select one)

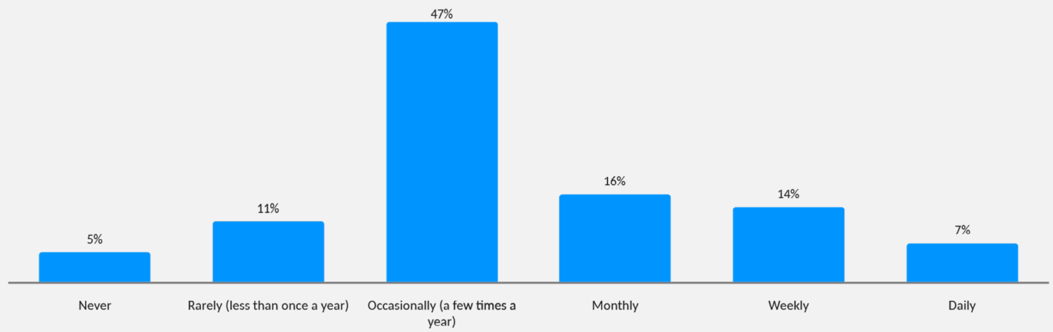

2025 CNCF Annual Survey, Q42, Sample size = 183, DKNS excluded

FIGURE 4

TYPES OF AI/ML WORKLOADS

Are you running AI/ML workloads on Kubernetes? If so, which types of tasks and workloads? (select all that apply)

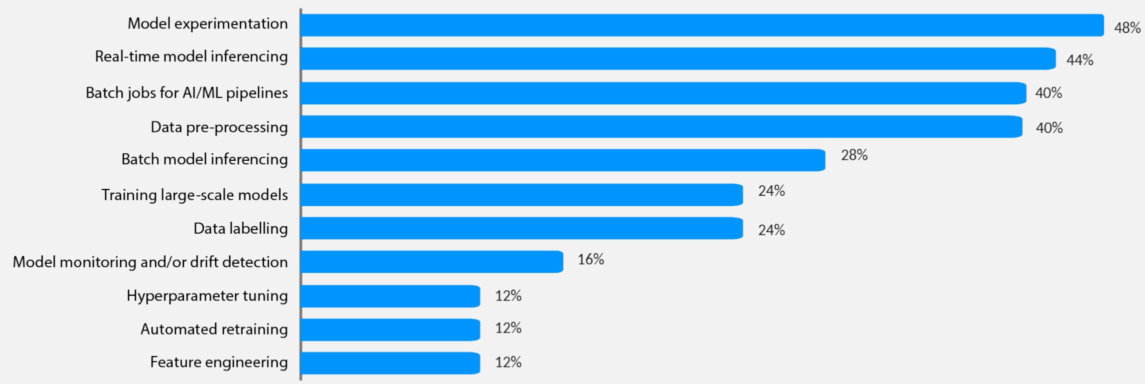

2025 CNCF Annual Survey Q29, Sample size = 81, shown to end-user organizations who use Kubernetes based on Q9, Q10 and Q23, DKNS and NA excluded

Infrastructure-first approach

The organizations succeeding with AI aren't just those with the best models. They're those with the infrastructure maturity to deploy and scale these workloads reliably. Kubernetes is emerging as the foundation, but success requires treating AI/ML as a first-class infrastructure challenge, not just an algorithmic one.

As organizations race to deploy AI workloads, a September 2025 open letter from open source infrastructure stewards issued a stark warning. Critical systems operate under "a dangerously fragile premise," relying on goodwill rather than sustainable funding models aligned with usage. For AI/ML workloads, which the letter explicitly identifies as drivers of "machine-driven, often wasteful automated usage", this sustainability challenge is acute.

The letter's authors note that "commercial-scale workloads often run without caching, throttling, or even awareness of the strain they impose". This is a description that fits AI workloads perfectly. Organizations deploying models occasionally (47% of respondents) may assume minimal infrastructure impact. However, outsized strain is still created on the infrastructure.

The path forward for AI requires an infrastructure-first approach. This means implementing caching strategies, using resource quotas, monitoring consumption, and contributing to the open source projects that make AI workloads possible. The CNCF ecosystem provides tools for sustainable orchestration, but only if organizations use them intentionally.

The maturity of cloud native infrastructure

Years of CNCF survey data tells us that cloud native adoption has matured from experimental to standard, Kubernetes has become infrastructure, and the biggest bottleneck may not be the technology but organizational change and the changing regulatory environment.

Between 2023 and 2025, the cloud native landscape shifted. What was once a forward-thinking architectural choice has become table stakes, where 98% of organizations now use cloud native techniques in at least some capacity, with early-stage adoption dropping to just 8%. Container usage in production applications increased from 41% to 56%, while Kubernetes cemented its position as the de facto orchestration platform, running in 82% of containerized environments.

Cloud native adoption is still high

Organizations reporting extensive use ("much" or "nearly all" development and deployment) grew from 54% in 2023 to 60% in 2024, and stayed at 59% in 2025 (Figure 5). Meanwhile, organizations just beginning or not yet using cloud native technologies dropped from 13% in 2023 to just 10% in 2025, suggesting the technology has moved well beyond early adoption.

FIGURE 5

CLOUD NATIVE ADOPTION

To what extent has your organization adopted cloud native techniques? (select one)

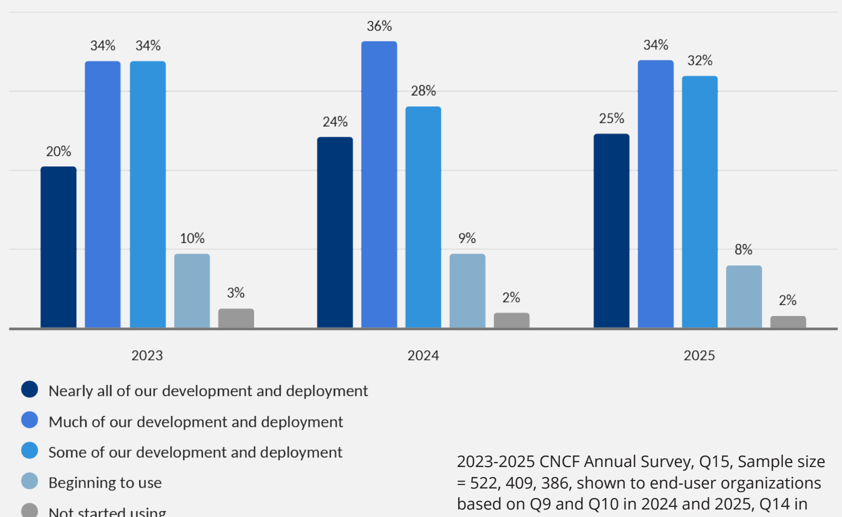

2023-2025 CNCF Annual Survey, Q15, Sample size = 522, 409, 386, shown to end-user organizations based on Q9 and Q10 in 2024 and 2025, Q14 in 2023, DKNS excluded

Container usage steadily increasing

The percentage using containers for most or all production applications increased from 41% in 2023 to 56% in 2025 (Figure 6). Organizations still piloting containers dropped from 11% to just 6%. This increase in two years represents significant progress, reflecting the maturity of container tooling where Docker and containerd provide reliable runtimes, image registries offer secure storage, and security scanning tools identify vulnerabilities.

Organizations have made decisions: either containers fit their needs and they move to production, or they don't and they abandon pilots. Few remain in indefinite experimentation.

As shown in Figure 7, the top challenge in 2025 is "cultural changes with the development team" (47%), followed by lack of training (36%) and security (36%). This marks a shift from 2023 when technical challenges like security and complexity dominated. Security will still be important for the years to come with new regulations coming in, such as the CRA.

Cultural resistance may manifest in multiple ways across organizations. Developers can express skepticism that containers add unnecessary complexity for simple applications or question whether Kubernetes is production-ready. Operations teams resist what they perceive as developer toys and express concern about troubleshooting containerized systems. Management worries about costly distractions from feature delivery and dependencies on specialized knowledge.

FIGURE 6

CONTAINER USAGE

How are containers used within your organization? (select one)

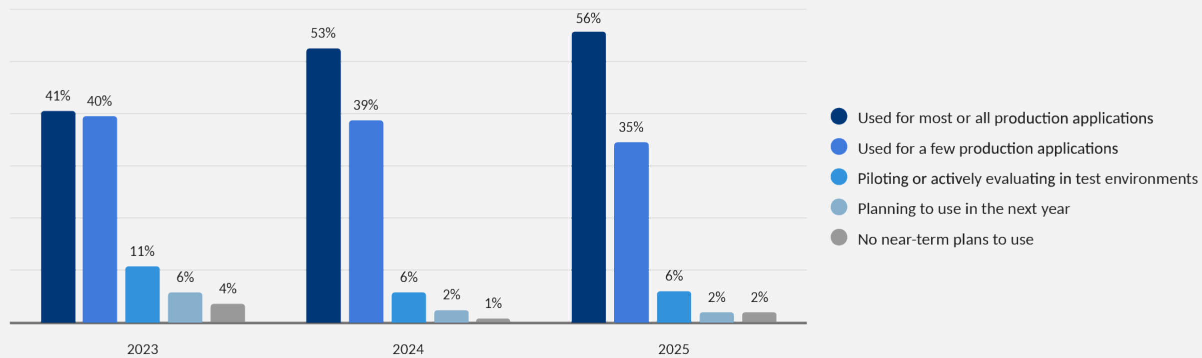

2023-2025 CNCF Annual Survey, Q19, Q20, Sample size = 522, 408, 365, shown to enduser organizations based on Q9 and Q10 in 2024 and 2025, Q14 in 2023, DKNS excludedexcluded

FIGURE 7

THE CHALLENGES OF CONTAINER USE

What are your challenges in using / deploying containers? (select all that apply)

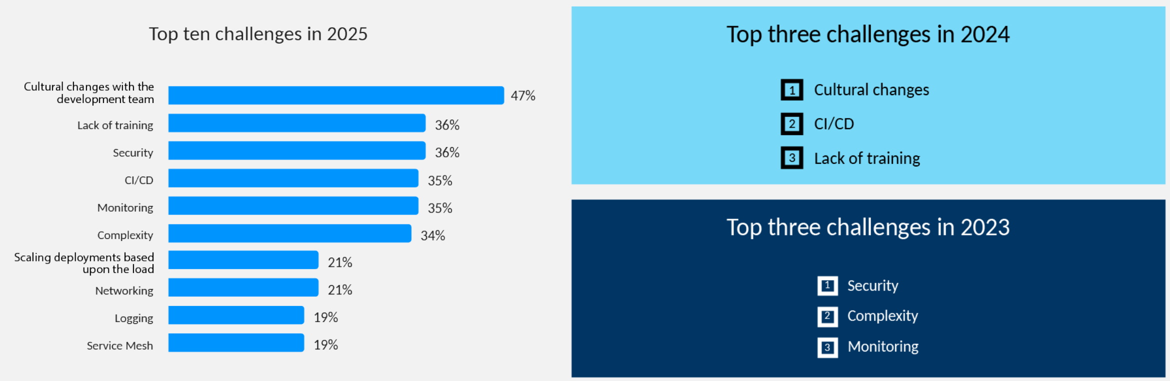

2023-2025 CNCF Annual Survey, Q21, Q22, Sample size = 477, 373, 351, Total mentions = 1,720, 1,425, 1,312, shown to end-user organizations based on Q9 and Q10 in 2024 and 2025, Q14 in 2023, DKNS excluded

“Kubernetes is boring” as the highest praise

Among container users, 82% are using Kubernetes in production in 2025, up from 66% in 2023 (Figure 8). This represents near-universal adoption within the container ecosystem. The characterization of Kubernetes as “boring” uses the term as highest praise. In technology, boring means reliable without unexpected failures, predictable with well-understood and documented behavior, mature with edge cases discovered and addressed, and stable with APIs that don’t break with every release.

This increase in two years reflects technology maturation where Kubernetes removed deprecated features and stabilized APIs, major cloud providers achieved feature parity, Helm charts and operators simplified application deployment, and the CRD ecosystem matured. Kubernetes won because it reached a standard stage and has network effects through ecosystem, knowledge, and tooling that alternatives can’t match.

FIGURE 8

KUBERNETES USAGE

Does your organization use Kubernetes? (select one)

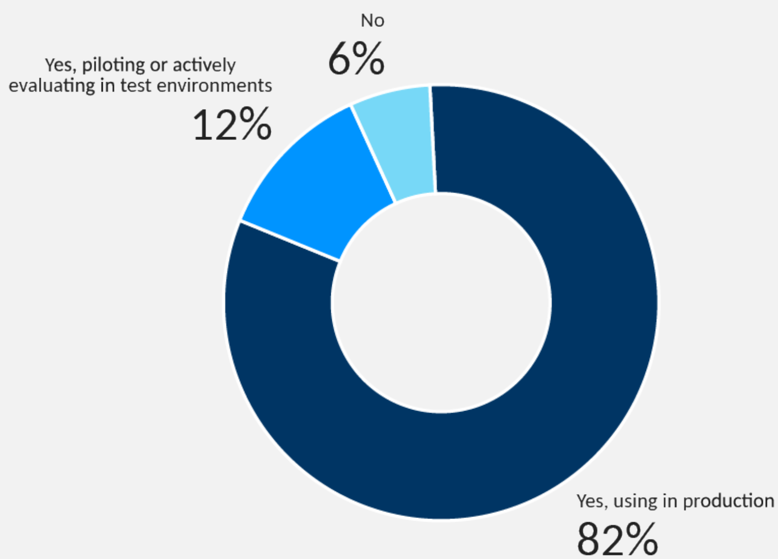

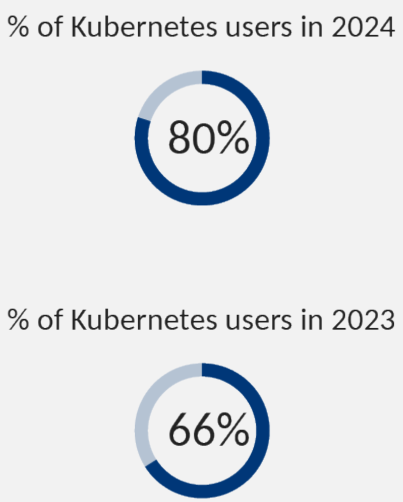

2023-2025 CNCF Annual Survey, Q22, Q23 Sample size = 477, 373, 351, shown to end-user organizations based on Q9 and Q10 in 2024 and 2025, Q14 in 2023, DKNS excluded.

WebAssembly still waiting for the inflection point

About 65% of organizations report no WebAssembly experience consistently across all three years, with only 5% having full deployment experience in 2025 (Figure 9). This suggests WebAssembly hasn't yet had its inflection point in cloud native environments. WebAssembly offers compelling advantages including language agnosticism, near-native performance, sandboxed security, portability, and small footprint. In theory, Wasm could replace containers for many workloads, offering faster cold starts, better density, and improved security.

FIGURE 9

WEBASSEMBLY ADOPTION

Have you or your organization ever deployed an application using WebAssembly? (select one)

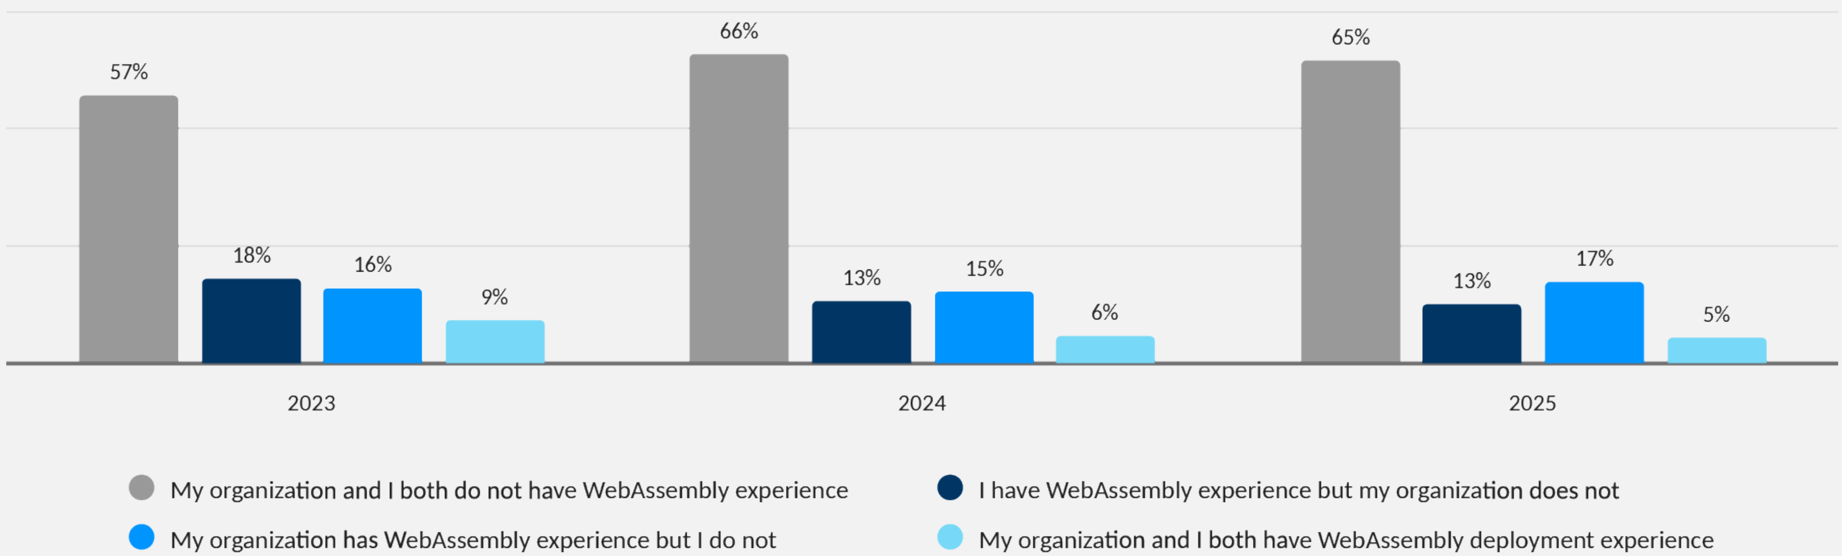

2023-2025 CNCF Annual Survey, Q37, Q39, Q43, Sample size = 988, 403, 347, DKNS excluded

Cloud native maturity profiles

To better understand where organizations are in their cloud native journey, we categorized them by their native adoption levels resulting in four distinct maturity profiles:

- **Cloud native explorers (8% of organizations):** Beginning to use cloud native techniques. These organizations experiment with containers and basic deployments.
- **Cloud native adopters (32%):** Using cloud native for some of their development and deployments. These organizations have selective application to specific projects or teams.
- **Cloud native practitioners (34%):** Using cloud native for much of their development and deployment. These organizations have mainstream adoption across most initiatives.
- **Cloud native innovators (25%):** Using cloud native for nearly all development and deployment. These organizations achieved comprehensive, organization-wide transformation.

The survey reveals a maturity progression model in cloud native adoption. It shows that as organizations advance from "explorers" to "innovators" (based on how extensively they use cloud native techniques), they systematically adopt more sophisticated practices and tooling. It's an illustration that cloud native maturity isn't just about running containers. It's also about embracing an entire ecosystem of modern development practices.

FIGURE 10

CLOUD NATIVE ADOPTION

To what extent has your organization adopted cloud native techniques? (select one)

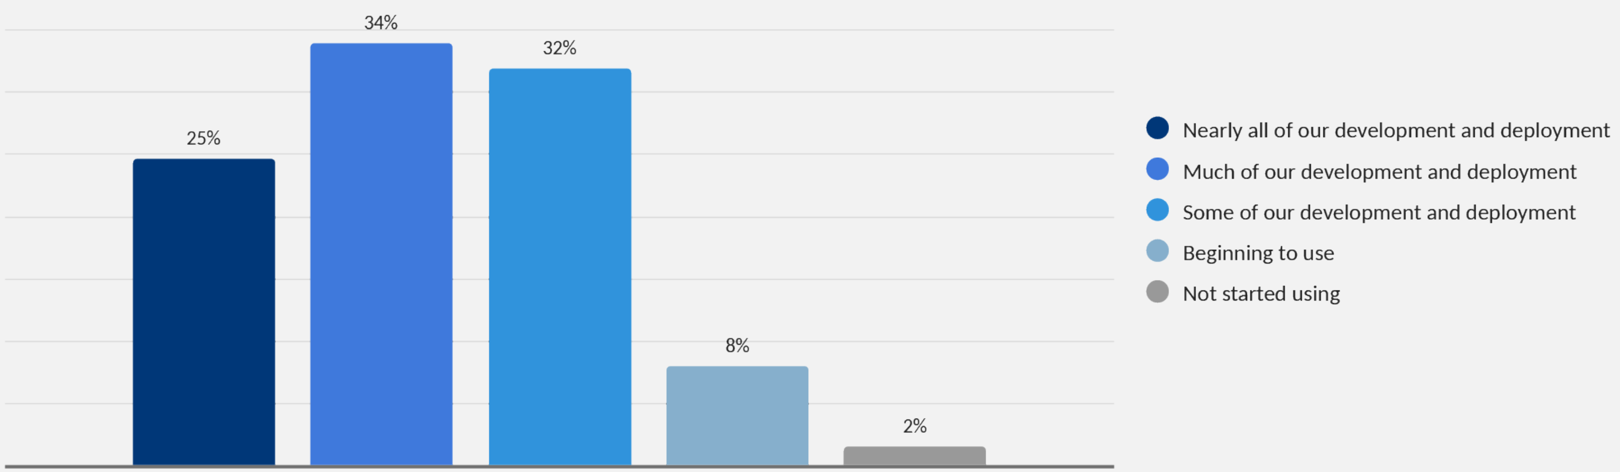

2025 CNCF Annual Survey, Q15, Sample size = 386, shown to end-user organizations based on Q9 and Q10, DKNS excluded

Cloud native explorers

Cloud native explorers are now the exception rather than the rule, comprising just 8% of organizations (Figure 10). As shown in Figure 11, The heavy representation of large enterprises (45% with 5,000+ employees) suggests that size creates complexity. The large enterprise skew reveals that thousands of applications across diverse technology stacks, decades of accumulated technical debt, risk-averse cultures requiring extensive evaluation, and distributed decision-making slow adoption for these organizations. These organizations are in learning mode, with minimal revenue tied to cloud native technologies. Averaging at 10%, indicating that cloud native is experimental rather than business-critical, with traditional systems still generating revenue.

Cloud native adopters

Adopters are in the early expansion phase. The average revenue they derive from cloud native is only 26%, suggesting these organizations are still building internal capabilities before further commercializing their expertise (Figure 12).

Geographic distribution shows Europe leading at 58%, Americas at 29%, and Asia-Pacific at 13%. The modest revenue percentage reflects that cloud native workloads remain in development and staging rather than production, internal platforms are being built but not yet revenue-generating, and migration is in progress while traditional systems remain dominant.

FIGURE 11

CLOUD NATIVE EXPLORERS

Organizations beginning to use cloud native computing techniques

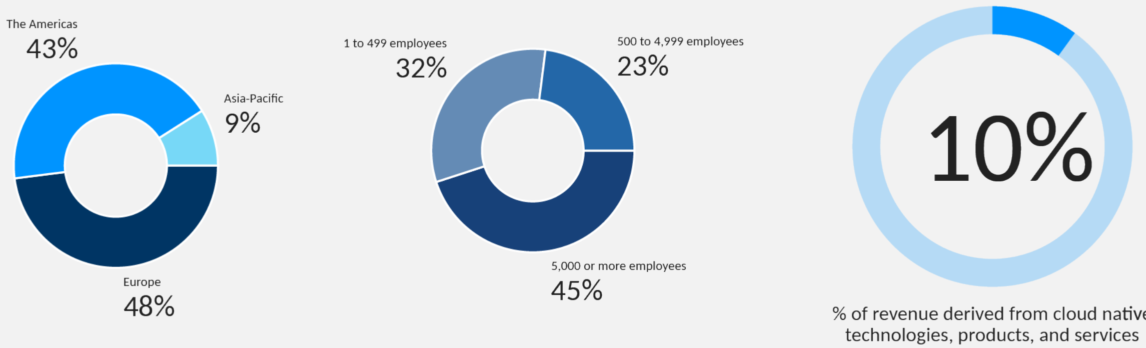

2025 CNCF Annual Survey, Q12, Q13, Q14 by Q15, Sample size = 31 (organizations beginning to use cloud native)

FIGURE 12

CLOUD NATIVE ADOPTERS

Organizations using cloud native computing techniques for some of their development

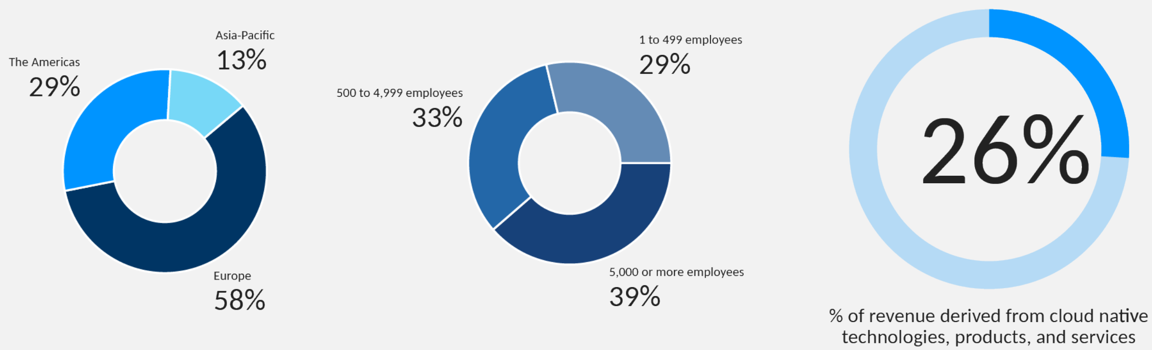

2025 CNCF Annual Survey, Q12, Q13, Q14 by Q15, Sample size = 123 (organizations using cloud native for some of their development)

Cloud native practitioners

Organizations at this maturity level typically treat cloud native as the default for new development, have platform engineering teams providing self-service capabilities, use GitOps workflows as standard across teams. As shown in Figure 13, the average revenue derived from cloud native technologies is 35%, indicating that advanced adoption correlates with business model evolution. This revenue inflection point suggests that the majority of production workloads run on cloud native infrastructure, new products are built cloud-native-first, and legacy migration is substantially complete for tier-1 applications. They also likely maintain comprehensive observability through metrics, logs, and traces, automate security policies via admission controllers, operate multi-cluster deployments across environments, and regularly test disaster recovery procedures.

Cloud native innovators

Small companies (1-499 employees) dominate the innovator category at 55%, suggesting agility advantages (Figure 14). Over half of their revenue is derived from cloud native technologies. This can show startup advantages where there may be no legacy infrastructure to migrate, technical cultures are engineering-driven with DevOps DNA, small teams can transform faster than large enterprises, and economic necessity requires infrastructure efficiency to compete. The majority revenue share indicates that cloud native is core infrastructure rather than experiment, business models depend on cloud native capabilities, and competitive advantage is derived from infrastructure maturity as well as cloud native innovation.

FIGURE 13

CLOUD NATIVE PRACTITIONERS

Organizations using cloud native computing techniques for much of their development

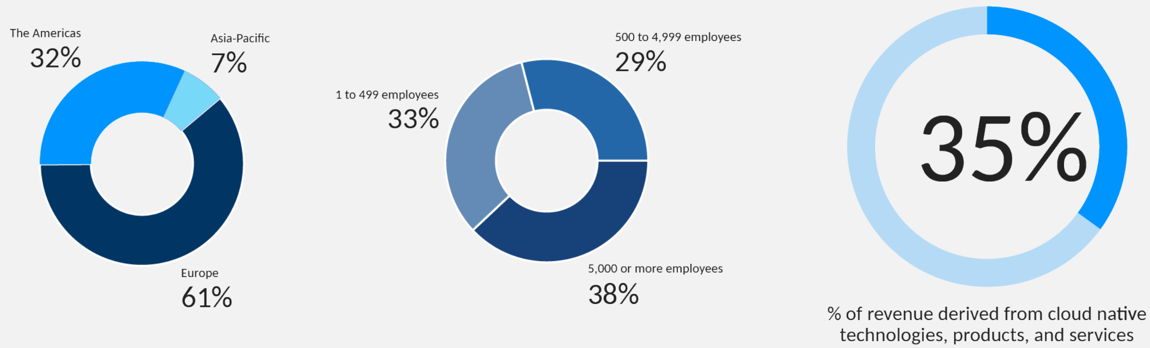

2025 CNCF Annual Survey, Q12, Q13, Q14 by Q15, Sample size = 131 (organizations using cloud native for much of their development)

FIGURE 14

CLOUD NATIVE INNOVATORS

Organizations using cloud native computing techniques for nearly all of their development

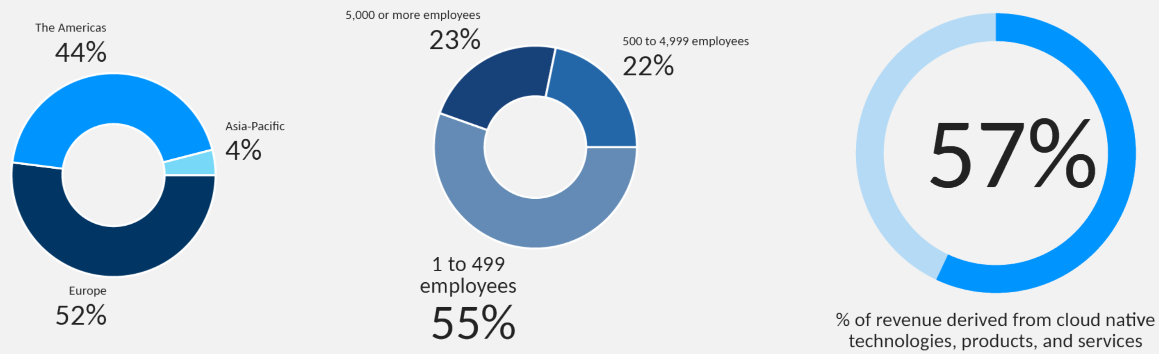

2025 CNCF Annual Survey, Q12, Q13, Q14 by Q15, Sample size = 95 (organizations using cloud native for nearly all oftheir development)

Technology roadmap and release practices

Advanced technologies require foundational maturity. Innovators are almost 3x more likely than explorers to run service mesh in production. And stateful containers and serverless show high adoption among mature organizations. Development velocity is a clear differentiator between maturity levels. Innovators operate at a fundamentally different cadence, where they check in code more often and automating deployments at rates explorers and adopters haven't yet approached.

The data reveals that as organizations advance from explorers to innovators, they systematically adopt more sophisticated practices and tooling (Figure 15). Technology adoption follows maturity, where more mature organizations have universally adopted core technologies such as stateful containers at 79% for innovators, serverless at 64%, and service mesh at 39%. Development velocity scales with maturity, where innovators check in code multiple times daily at 74%, run daily releases at 41%, and automate most deployments at 59%, practices that explorers barely engage in (Figure 16).

FIGURE 15

TECHNOLOGY ROADMAP

% of organizations running in production

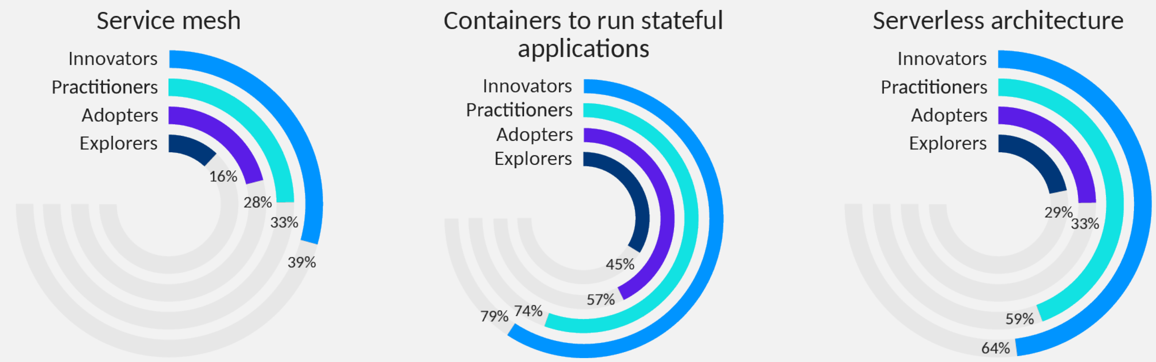

2025 CNCF Annual Survey, Q51-53 by Q15, Sample size = 380, Q15 was only shown to end-user organizations

FIGURE 16

CODE, RELEASE CYCLES, AND AUTOMATION

% of organizations

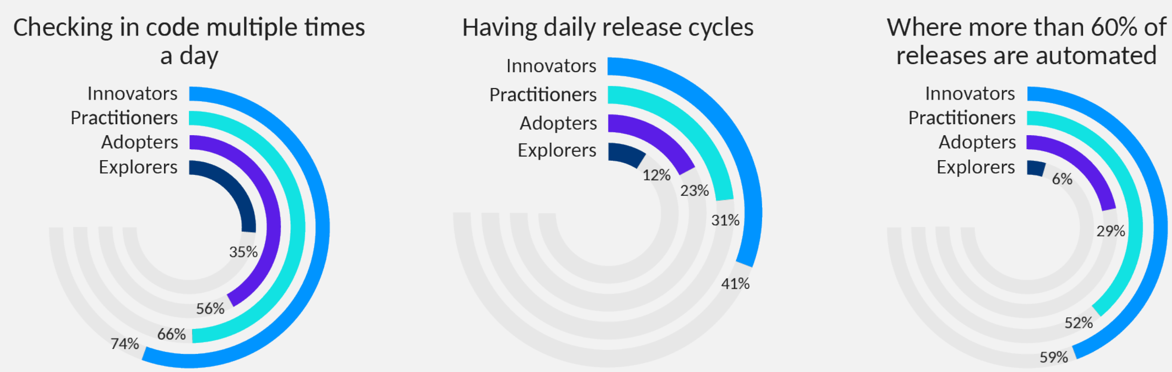

2025 CNCF Annual Survey, Q56-58 by Q15, Sample size = 380, Q15 was only shown to end-user organizations

GitOps and CI/CD

As Figure 17 shows,GitOps represents advanced cloud native maturity. Zero percent of explorers have embraced it, while 58% of innovators run GitOps-compliant deployments. GitOps is a capstone practice that requires extensive foundational work.

CI/CD is the gateway to cloud native maturity (Figure 18). Even organizations just starting their journey have adopted it at 42%, and it becomes nearly universal (91%) among innovators, making it perhaps the most fundamental practice in the cloud native toolkit.

FIGURE 17

GITOPS PRINCIPLES

% of organizations where much or all of their deployment practices and tools adhere to GitOps principles

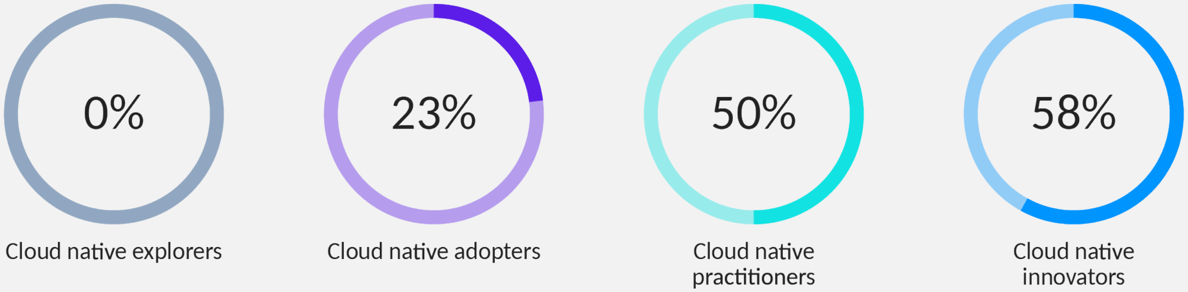

2025 CNCF Annual Survey, Q59 by Q15, Sample size = 380, Q15 was only shown to end-user organizations

FIGURE 18

CI/CD TOOLS

% of organizations using CI/CD tools in production to manage pipelines

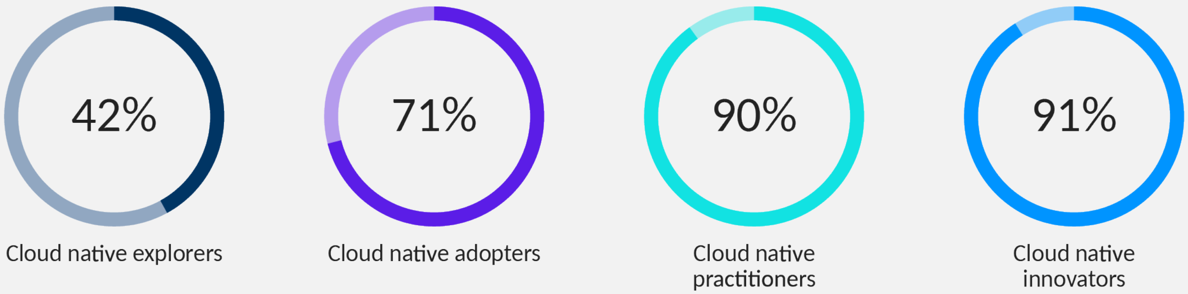

2025 CNCF Annual Survey, Q60 by Q15, Sample size = 380, Q15 was only shown to end-user organizations

CNCF Projects

Graduated projects

CNCF graduated projects represent the most mature, production-ready technologies in the cloud native ecosystem (Figure 19). These projects have demonstrated widespread adoption, strong governance, and sustained contribution patterns. The data reveals a hierarchy of adoption with foundational infrastructure projects like Kubernetes, Helm, and etcd achieving near-universal production use (81-87%). Newer or more specialized projects show lower but growing adoption rates. The 'evaluating' percentages indicate ongoing interest and future growth potential, particularly for emerging projects like KEDA (16% evaluating), Open Policy Agent (20%) and Argo (20%).

Core observability and security projects show strong but varied adoption. Prometheus reaches 77% production use with 12% evaluating, CoreDNS achieves 76% with 5% evaluating, and containerd reaches 74% with 7% evaluating. These tools have become standard components of cloud native infrastructure. Cert Manager at 58% production use and Argo at 52% demonstrate solid adoption for specific use cases around certificate management and GitOps workflows. Projects such as fluentd at 41%, Istio at 36%, cri-o at 34%, Envoy at 33%, and Harbor at 32% show more specialized adoption patterns reflecting their use in specific architectural patterns or as alternatives within their respective categories.

FIGURE 19

CNCF GRADUATED PROJECTS

Which of these graduated CNCF projects is your organization using in production or evaluating? (select one response)

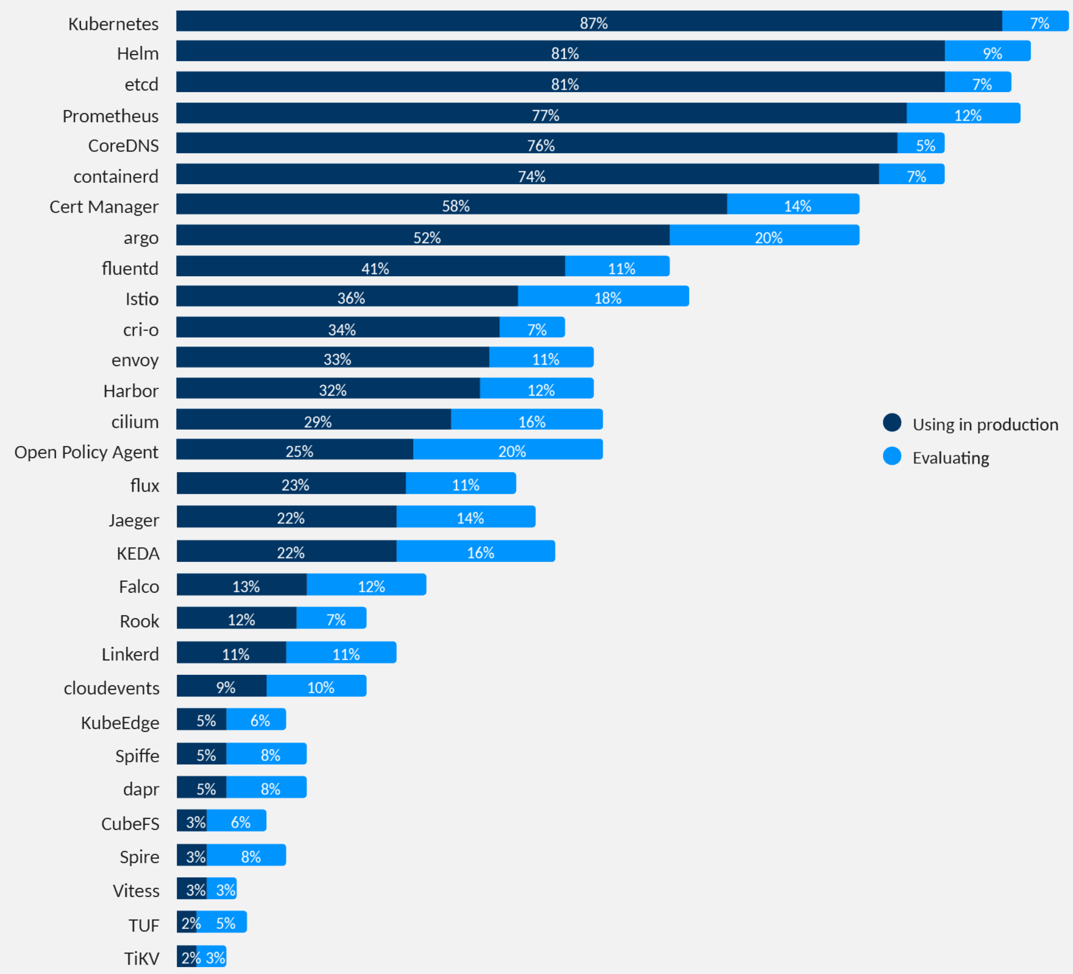

2025 CNCF Annual Survey, Q32, Sample size = 395-596, DKNS excluded

Incubating projects

CNCF incubating projects represent the next generation of cloud native infrastructure technologies that have proven viability but haven't yet achieved the widespread adoption of graduated projects (Figure 20). These projects address emerging needs, fill ecosystem gaps, and often target specific use cases rather than universal problems. The data reveals a wide spectrum: from nearly-graduated projects with 50%+ production use to early-stage incubations with single-digit adoption.

Leading incubating projects include CNI at 52% production use, demonstrating its critical role in Kubernetes networking. OpenTelemetry reaches 49% production use with 26% evaluating, showing strong momentum in the observability space as organizations standardize on this emerging standard. gRPC at 44% and KeyCloak at 42% show solid adoption for their respective domains of remote procedure calls and identity management. Projects such as Operator Framework, Thanos, and Crossplane demonstrate growing adoption as organizations mature their Kubernetes operations and multi-cloud strategies.

Emerging projects show varied adoption reflecting their specialized nature or earlier stage of development. Chaos Mesh at 11% evaluating alongside 3% production use indicates growing interest in chaos engineering practices. Projects like KubeVirt, Backstage, and Keptn show single-digit production use but healthy evaluation numbers, suggesting they're gaining traction in their specific domains of virtual machine management, developer portals, and application lifecycle orchestration respectively.

FIGURE 20

CNCF INCUBATING PROJECTS

Which of these incubating CNCF projects is your organization using in production or evaluating? (select one)

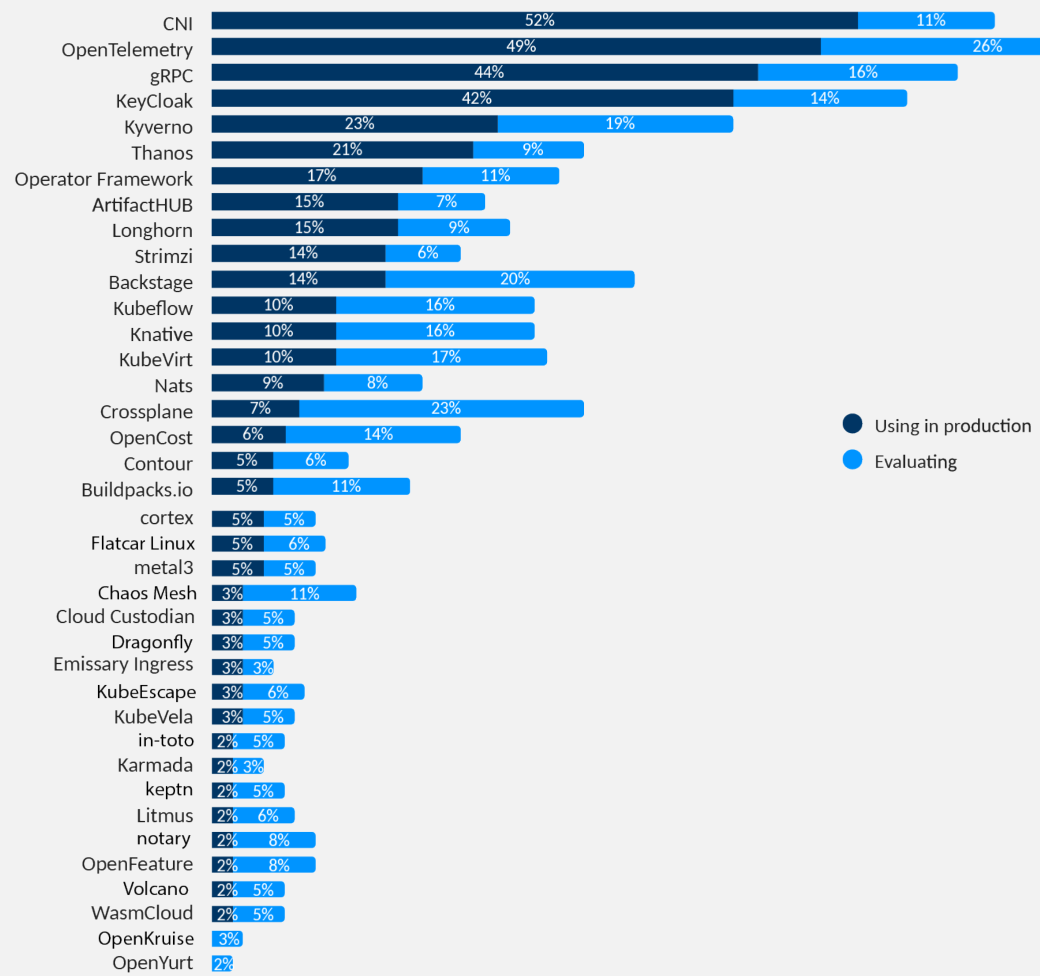

2025 CNCF Annual Survey, Q33, Sample size = 395-502, DKNS excluded

Conclusion

With 98% adoption across organizations, the conversation has shifted from whether to adopt cloud native to how to maximize its value. The convergence of AI workloads and cloud native infrastructure represents the most significant development of 2025, with 66% of organizations running generative AI on Kubernetes. However, the gap between AI ambition and deployment reality remains stark. While organizations rush to experiment with AI models, only 7% achieve daily deployments, and most consume rather than train models.

The progression from explorer to innovator follows predictable patterns, with GitOps adoption serving as a reliable indicator of organizational maturity. Organizations at 0% GitOps adoption remain in early stages, while those at 58% have achieved comprehensive transformation. This predictability provides a roadmap for organizations seeking to advance their cloud native journey.

The sustainability of open source infrastructure has emerged as an important concern. As AI workloads strain systems through machine-driven automated usage, the open source projects underpinning this infrastructure face unprecedented pressure. Organizations benefiting from cloud native must move beyond passive consumption to active stewardship through funding, contribution, and responsible resource usage. The alternative is a tragedy of the commons where critical infrastructure degrades under load.

Looking forward, cloud native is no longer the destination but the foundation. The organizations succeeding in 2025 and beyond will be those that treat infrastructure as a first-class capability, invest in organizational transformation alongside technical adoption, and recognize that sustainable infrastructure requires sustainable funding models. The survey data provides not just a snapshot of current adoption but a preview of what separates leaders from followers in an increasingly infrastructure-dependent world.
As cloud native becomes boring infrastructure, the competitive advantage shifts to those who can build reliable, scalable, and sustainable systems on that foundation.

Methodology

This study is based on a web survey conducted by Linux Foundation Research and its partners during September 2025. The survey's goal was to understand organizational perspectives on containers, Kubernetes, and other cloud native adoption trends. In this section, we present the study methodology and context regarding how we analyzed the data followed by the demographics of the respondents.

From a research perspective, it was important to ensure high data quality. We addressed data quality through extensive prescreening, survey screening questions, and data quality checks to ensure that respondents had sufficient professional experience to answer questions accurately on behalf of the organization they worked for.

We collected survey data from industry-specific companies, IT vendors and service providers, nonprofit, academic, and government organizations. Respondents spanned many vertical industries and companies of all sizes, and we collected data from several geographies, including the Americas, Europe, Asia-Pacific, and the Middle East and Africa.

The 2025 CNCF Annual Survey comprised 65 questions that addressed screening, respondent demographics, cloud native computing, containers, Kubernetes, CNCF projects, and six other topical areas. The high-level design of the survey is shown in Table 1.

TABLE 1

SURVEY DESIGN

Section	Questions	Question categories	Who answers the questions
Demographics	Q1 – Q14	Tell us about yourself & the organization you work for	All respondents except unemployed students
Cloud Native	Q15 – Q19	Cloud and cloud native computing at your organization	For orgs who are end-users, SIs or consultants
Containers	Q20 – Q22	Container use	For orgs who are end-users, SIs or consultants
Kubernetes	Q23 – Q31	Kubernetes use, follow-up, & autoscaling	For orgs who are end-users, SIs or consultants
CNCF Projects	Q32 – Q34	CNCF projects	All respondents
Security	Q35 – Q36	Security and compliance	All respondents
Observability	Q37 – Q38	Observability	All respondents
Generative AI	Q39 – Q42	Generative AI, self-hosting, Kubernetes use	All respondents
WebAssembly	Q43 – Q49	WebAssembly	All respondents
Tech Roadmap	Q50 – Q53	Technology roadmap	All respondents
Infrastructure	Q54 – Q55	Infrastructure platforms	All respondents
CI/CD	Q56 – Q61	Code, release cycles, and automation	All respondents
Students only	Q62 – Q65	Student only technology questions	Students only

Survey screening involves the use of three variables to select respondents who have sufficient experience that will allow them to answer some or all questions in the survey.

- Must be familiar with cloud native technologies
- Must identify as a human
- Must be employed full-time or part-time

A total of 628 respondents completed the survey. The margin of error for this sample size was +/- 3.3% at a 90% confidence level. The data was primarily segmented by geographic region, company size, type of organization, and extent of organizational adoption of cloud native techniques.

Although respondents were required to answer nearly all questions in the survey, a provision was made when a respondent was unable to answer a question. This is accomplished by adding a "Don't know or not sure" (DKNS) response to the list of responses for every question. However, this creates a variety of analytical challenges.

One approach was to treat a DKNS just like any other response so that the percentage of respondents that answered the DKNS is known. The advantage of this approach is that it shows the exact distribution of data collected. The challenge with this approach is that it can distort the distribution of valid responses, i.e., responses where respondents could answer the question.

Some of the analyses in this report exclude DKNS responses. This is done when comparing data year-over-year because the sample size and percentage of DKNS responses always varies from year to year. Excluding DKNS responses enables us to normalize the data so that year-over-year comparisons and growth rates can be accurately calculated.

The percentage values in this report may not total to exactly 100% due to rounding.

Availability of results

LF Research makes each of its empirical project datasets and this 2025 CNCF Annual Survey available on Data.World. Included in this dataset are the survey instrument, raw survey data, screening and filtering criteria, and frequency charts for each question in the survey. LF research datasets, including this project, can be found at data.world/thelinuxfoundation. Access to Linux Foundation datasets is free but does require you to create a data.world account.

Respondent demographics

These demographics provide you with a profile of the 2025 CNCF Annual Survey respondents (Figures 21, 22 and 23). For the original source data and study frequencies, please see the data.world access described above.

FIGURE 21

GEOGRAPHIC DISTRIBUTION OF SURVEYED ORGANIZATIONS

What is the geographic location of your organization’s headquarters? (select one)

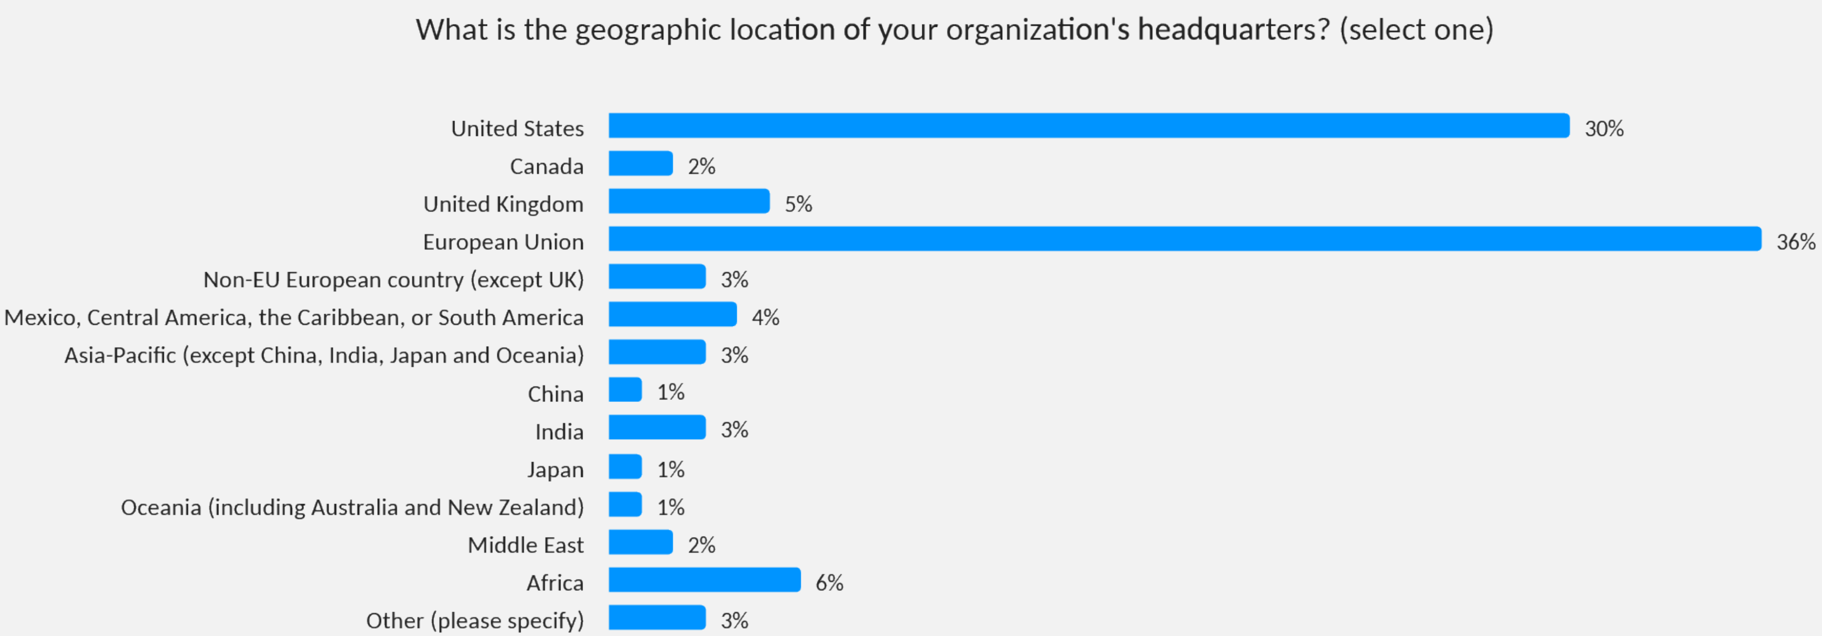

2025 CNCF Annual Survey, Q12, Sample size = 628

FIGURE 22

SELECTED DEMOGRAPHICS I

What title most accurately describes your role? (select one)

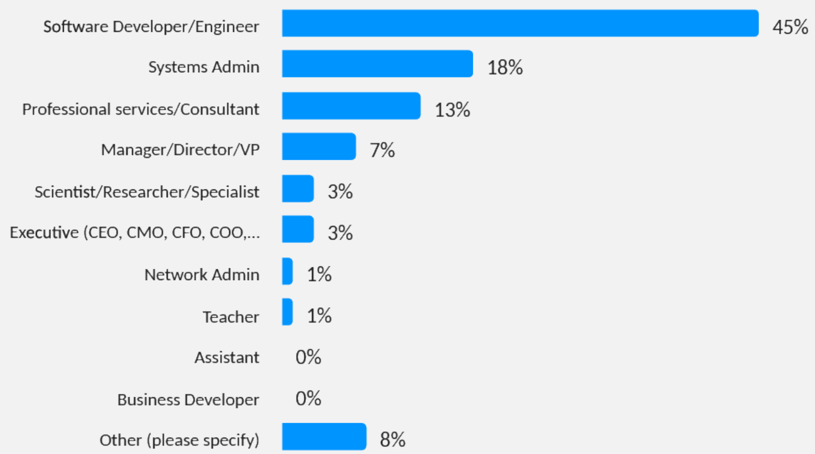

How many employees does your organization have worldwide?

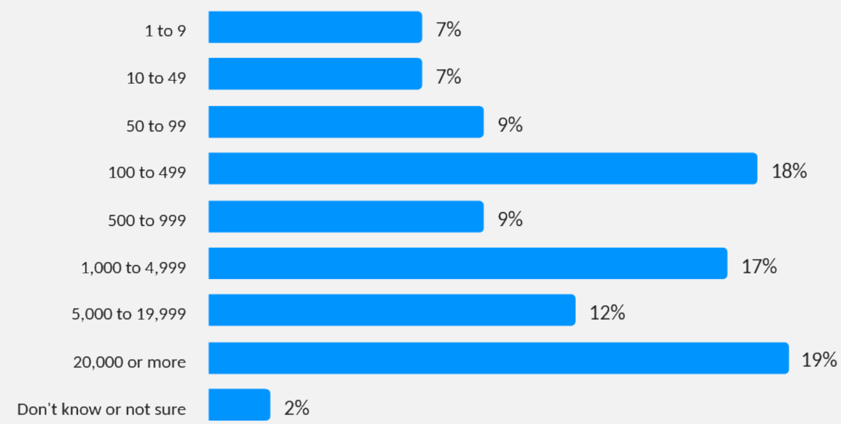

2025 CNCF Annual Survey, Q8, Q13, Sample size = 628

FIGURE 23

SELECTED DEMOGRAPHICS II

How many years of professional experience in IT do you have? (select one)

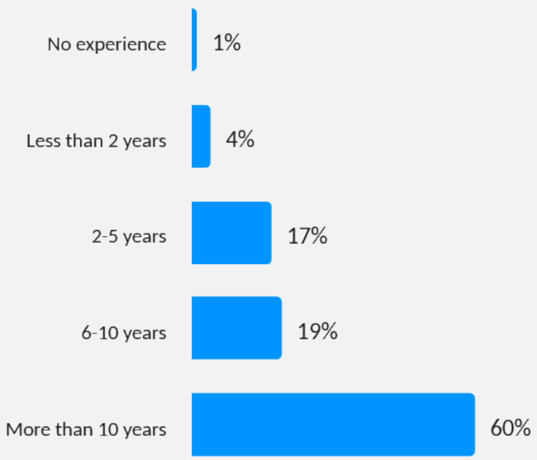

Which option best describes the organization you work for? (select one)

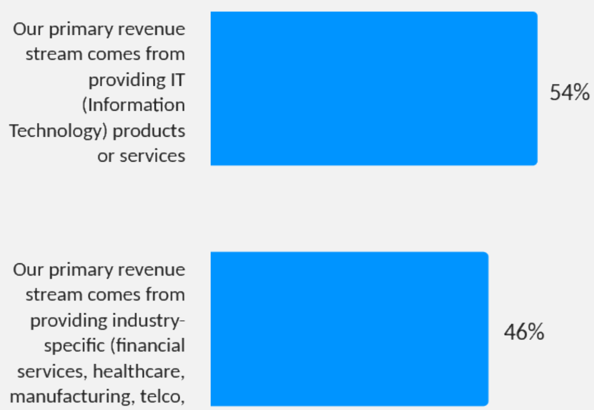

About how much of your organization's main source of revenue is derived from providing cloud native technologies, products, or services? (select one)

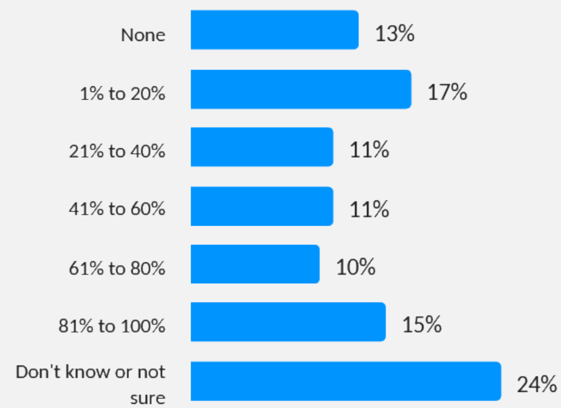

2025 CNCF Annual Survey, Q5, Q9, Q14, Sample size = 628

About the authors

Adrienn Lawson is Director of Quantitative Research at the Linux Foundation, where she leads data-driven initiatives to understand open source ecosystems. With expertise in social data science from the University of Oxford and a background spanning academic and governmental research, she brings methodological rigor to analyzing distributed collaboration networks. At the Linux Foundation, Adrienn leads a team conducting cross-sectional research across industry verticals and geographic regions to provide comprehensive insights into open source dynamics. Her work encompasses empirical investigations into regulatory compliance, the implications of AI, and sustainable funding models. She produces evidence-based recommendations that inform strategic decision-making within the open source community.

Jeffrey Sica is Head of Projects at the CNCF, with a focus on improving maintainer experience, building communities, and project automation. Before that, he worked at Red Hat and the University of Michigan focusing on cloud native technologies and CICD patterns. Jeffrey has been a contributor to upstream Kubernetes, helping in SIG-Contribex, SIG-Release, and SIG-UI. He passionately advocates for open source development and recognizing and alleviating burnout.

Acknowledgments

We thank all the people who participated in the survey. Special thanks to Linux Foundation colleagues for their involvement in the various stages of the research process: Chris Aniszczyk, Jonathan Bryce, Hilary Carter, Mia Chaszeyka, Anna Hermansen, and Allison Price as well as the Linux Foundation's Creative Services team.

Cloud native computing empowers organizations to build and run scalable applications with an open source software stack in public, private, and hybrid clouds. The Cloud Native Computing Foundation (CNCF) hosts critical components of the global technology infrastructure, including Kubernetes, Prometheus, and Envoy. CNCF brings together the industry's top developers, end users, and vendors and runs the largest open source developer conferences in the world. Supported by nearly 800 members, including the world's largest cloud computing and software companies, as well as over 200 innovative startups, CNCF is part of the nonprofit Linux Foundation. For more information, please visit www.cncf.io.

- https://twitter.com/cloudnativefdn
- https://www.facebook.com/CloudNativeComputingFoundation/
- https://www.linkedin.com/company/cloud-native-computing-foundation/
- https://www.youtube.com/c/cloudnativefdn
- https://github.com/cncf

Founded in 2021, Linux Foundation Research explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. By leveraging project databases and networks and committing to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations worldwide.

- twitter.com/linuxfoundation
- facebook.com/TheLinuxFoundation
- linkedin.com/company/the-linux-foundation
- youtube.com/user/TheLinuxFoundation
- github.com/LF-Engineering

To reference this work, please cite as follows: Adrienn Lawson and Jeffrey Sica, "CNCF Annual Cloud Native Survey: The infrastructure of AI's future" foreword by Jonathan Bryce, The Linux Foundation, January 2026.