RAFAY

2024 REPORT

The Pulse of Enterprise Platform Teams: Cloud, Kubernetes and AI

Table of Contents

SECTION 1

Executive Summary

Platform engineering teams have emerged as key drivers of digital transformation, influencing critical decisions on cloud infrastructure, Kubernetes and AI technologies. However, these teams face mounting challenges in managing costs and complexity across these domains. To succeed, organizations must equip their platform teams with advanced tools and strategies that optimize resources, standardize processes and foster innovation, ultimately positioning themselves for success in a technology-driven business environment.

Enterprise Platform Teams Expect Major Challenges in Kubernetes Scaling, Cloud Costs and Developer Productivity

Nearly all organizations reported that they face challenges in Kubernetes management.

More than 9 in 10

organizations report hurdles in Kubernetes management with their current solutions, as modern enterprises grapple with using a multitude of applications for their projects.

While Kubernetes and cloud infrastructures provide many benefits, cost visibility and control is a top challenge.

of organizations cite managing cost visibility and controlling Kubernetes and cloud infrastructure costs as a top current Kubernetes-related challenge.

Platform teams are stretched thin managing complex Kubernetes cluster lifecycles and cloud environments.

In addition to managing cost visibility and controlling cloud and Kubernetes costs, other top Kubernetes challenges include:

38%

The complexity of keeping up with Kubernetes cluster lifecycle management with multiple, disparate tools.

38%

The establishment and upkeep of enterprise-wide standardization.

To address these challenges, over the next year, cost management, automation and developer productivity are top platform team priorities.

report that reducing and optimizing costs associated with Kubernetes infrastructure remains a top cost management initiative.

Nearly half of organizations are prioritizing providing self-service experiences for developers as a future developer- and automation-focused initiative for the next year.

AI and GenAI: A New Frontier Mirroring Kubernetes Adoption Challenges

Organizations universally recognize the strategic importance of efficiency tools for AI and GenAI technologies.

96%

cite it is important for organizations to provide efficient methods for the development and deployment of AI applications

94%

say the same for GenAI applications

However, a majority are in the early stage of adoption and face experimentation and implementation hurdles.

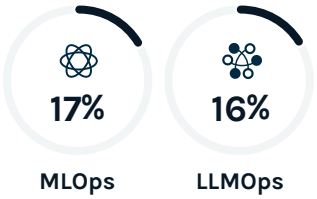

Less than a quarter of organizations are at a sufficient level of implementation for both MLOps and LLMOps

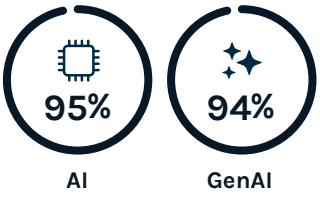

95% of teams with MLOps implementations report challenges in experimenting with and deploying AI apps, while 94% struggle with GenAI app experimentation and deployment

Platform teams are emerging as key facilitators of AI and GenAI implementation and governance, mirroring their responsibilities in the cloud and Kubernetes management space.

They will play a pivotal role in helping teams to overcome challenges with the development of AI and GenAI applications.

The top responsibilities organizations believe should be assigned to platform teams to assist in the development of AI/GenAI applications in their company include:

- **Security for MLOps and LLMOps workflows**
- **Model deployment automation**
- **Data pipeline management**

AI Surge and Kubernetes Expansion Drive Demand for Self-Service and Automation in Data Scientist and Developer Workflows

Organizations are prioritizing the developer experience with a growing emphasis on automation and self-service, spanning both AI initiatives and Kubernetes deployments.

Top developer- and automation-focused initiatives for the next year:

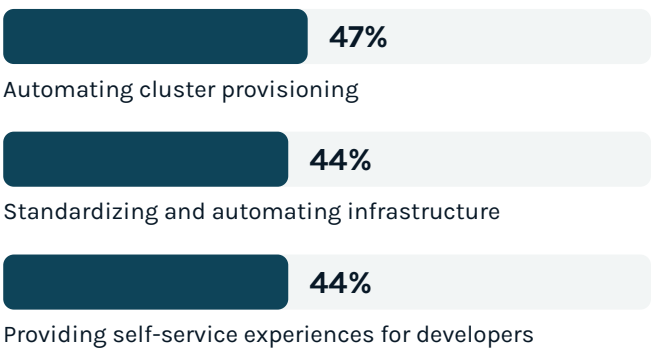

Through improved developer experiences in AI workflows, organizations expect substantial productivity improvements.

believe pre-configured AI workspaces with built-in MLOps and LLMOps tooling could save teams over 10% of time monthly.

In an organization with 100 developers earning an average salary of $140,000, this could unlock nearly 20,000 hours of developer time annually. This is equivalent to $1.4 million in salary costs or the productivity gain of 9 additional full-time developers — without increasing headcount.

With the right tools, organizations will also be well-positioned to empower platform teams to tackle critical Kubernetes infrastructure management initiatives in the year ahead.

Top platform team infrastructure priorities for the next year:

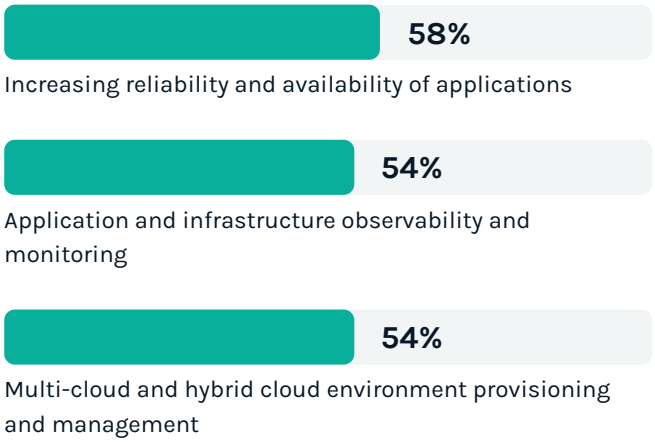

About This Study

This comprehensive study was conducted in two parts. The initial research gathered insights to explore the current and future trends and challenges enterprises and platform teams face in regards to Kubernetes and Kubernetes automation. It surveyed more than 1,000 platform engineering, platform architecture, cloud architecture, cloud engineering, developer, DevOps, site reliability engineering (SRE) and operations professionals with roles ranging from C-level to team members at U.S. organizations with over 1,000 employees.

Research was conducted to gain a better understanding of:

- The role of Kubernetes in enterprises and how organizations currently — and plan to — use it in the next year (e.g., challenges, opportunities, usage growth)
- The effects of Kubernetes challenges and cloud costs on developer and operations productivity
- The challenges platform teams face when provisioning environments and how to address them
- The initiatives organizations are prioritizing — with a focus on cost control and automation — so platform teams can better support developers

To complement this main survey, we conducted an additional study focused specifically on Artificial Intelligence (AI) and Generative AI (GenAI) adoption in the enterprise. This supplementary survey collected responses from 1,035 professionals mirroring the same industries and roles in the main survey.

This research provides valuable insights into:

- The current state of AI and GenAI initiatives
- The challenges organizations face
- The evolving role of platform teams in supporting these technologies

The findings from both surveys align closely, highlighting the interconnected nature of cloud, Kubernetes and AI technologies in modern enterprise environments.

SECTION 2

Current Trends and Challenges in Cloud and Kubernetes Automation

With the ongoing surge in Kubernetes usage, organizations are increasingly focused on managing the rising cost of Kubernetes and navigating the complexities of standardization. The cost and challenges of hybrid and multi-cloud management is also a concern and organizations are pointing to automation as a solution. Meanwhile, developers, operations and platform teams must sustain productivity amidst the integration of new technologies such as GenAI.

Surging Kubernetes Adoption Poses New Enterprise Challenges and Costly Price Tags

While **over one-third (35%)** of respondents are experienced with Kubernetes and have many clusters running in production, **nearly two-thirds (65%)** are still testing Kubernetes in a lab and training employees or only running a small number of clusters in production.

Challenges remain prevalent across organizations — **almost all respondents (93%)** report that their organizations encounter hurdles in Kubernetes management with their current solutions, as modern enterprises grapple with using a multitude of applications for their projects. They have also been inundated with a significant uptick in the cost and resources required to manage Kubernetes clusters.

The top challenges organizations face with Kubernetes are managing costs, handling cluster lifecycle complexity with multiple tools and maintaining enterprise-wide standardization. Cost visibility, infrastructure expense control and the intricacy of Kubernetes management emerge as key pain points for organizations.

Top Kubernetes Challenges

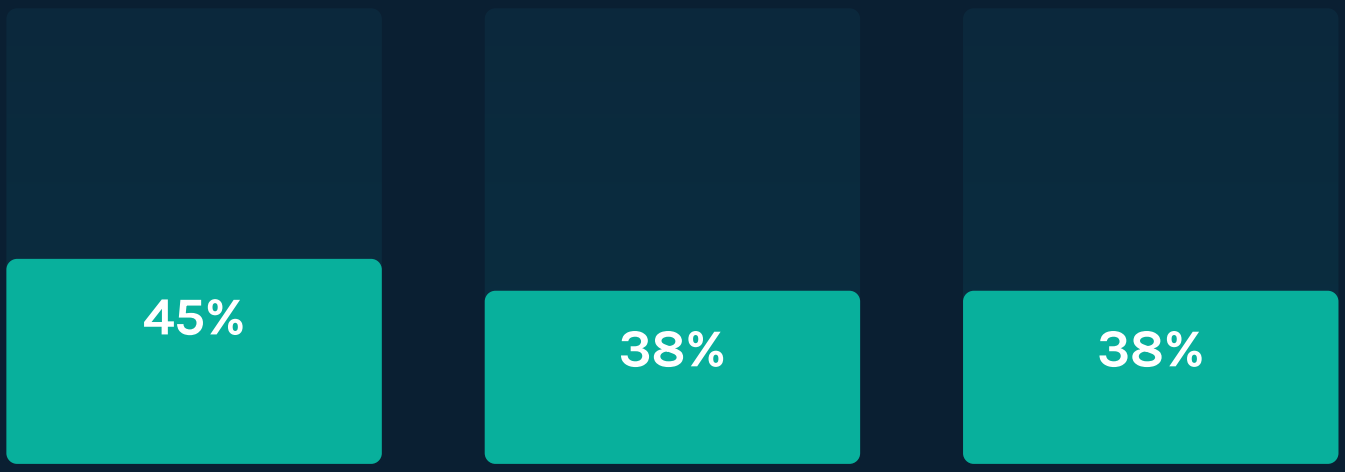

Managing cost visibility and controlling Kubernetes and cloud infrastructure costs

The complexity of keeping up with Kubernetes cluster lifecycle management with multiple, disparate tools

The establishment and upkeep of enterprise-wide standardization

Additional Kubernetes Challenges

In the coming year, organizations plan to tackle Kubernetes challenges by implementing solutions for observability and monitoring, improving cost visibility and enhancing security processes. These efforts will focus on streamlining application and infrastructure management while addressing key concerns in monitoring, cost control and security.

What does your organization plan to do to address these challenges with Kubernetes over the next year?

(Select all that apply)

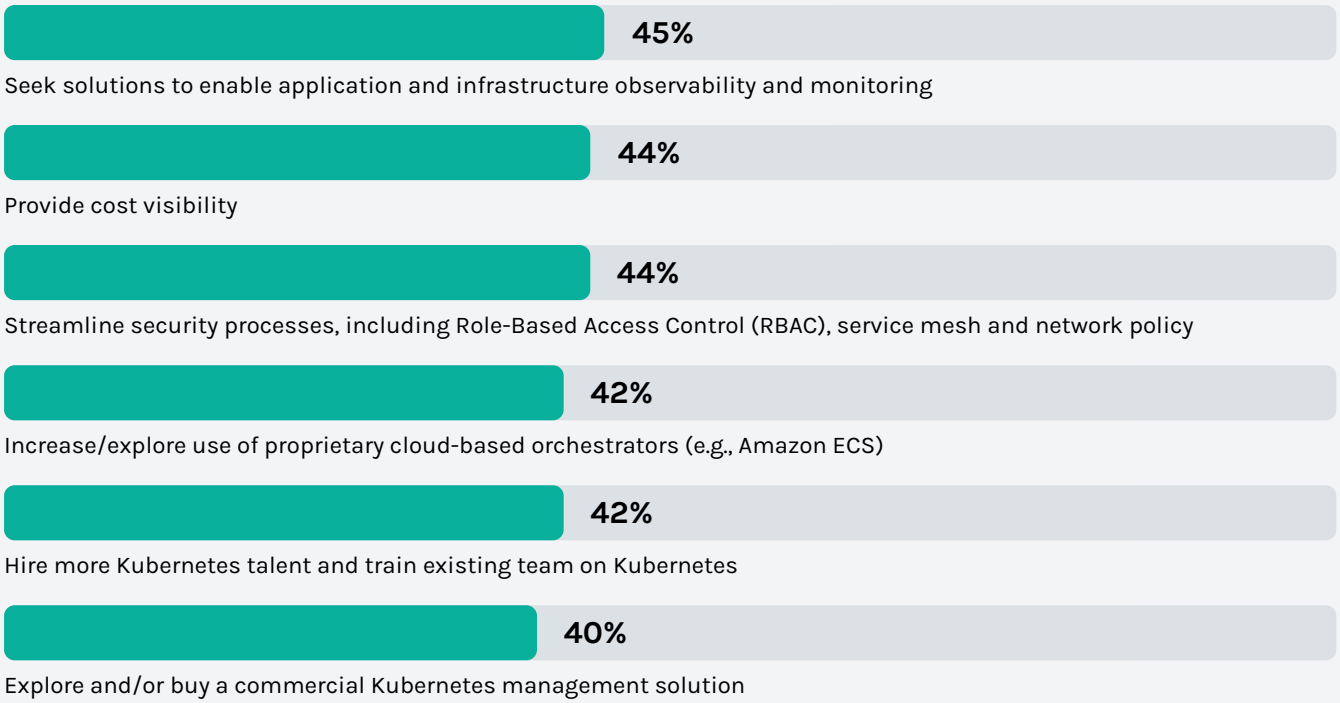

Strategies for Navigating Kubernetes Cost Challenges

Kubernetes is proving to be a costly pursuit. **Nearly one-third (31%)** state that the total cost of ownership (including software/support licenses, salaries of resources) is higher than budgeted for/anticipated.

In an effort to reduce and optimize these costs, organizations are applying various tactics.

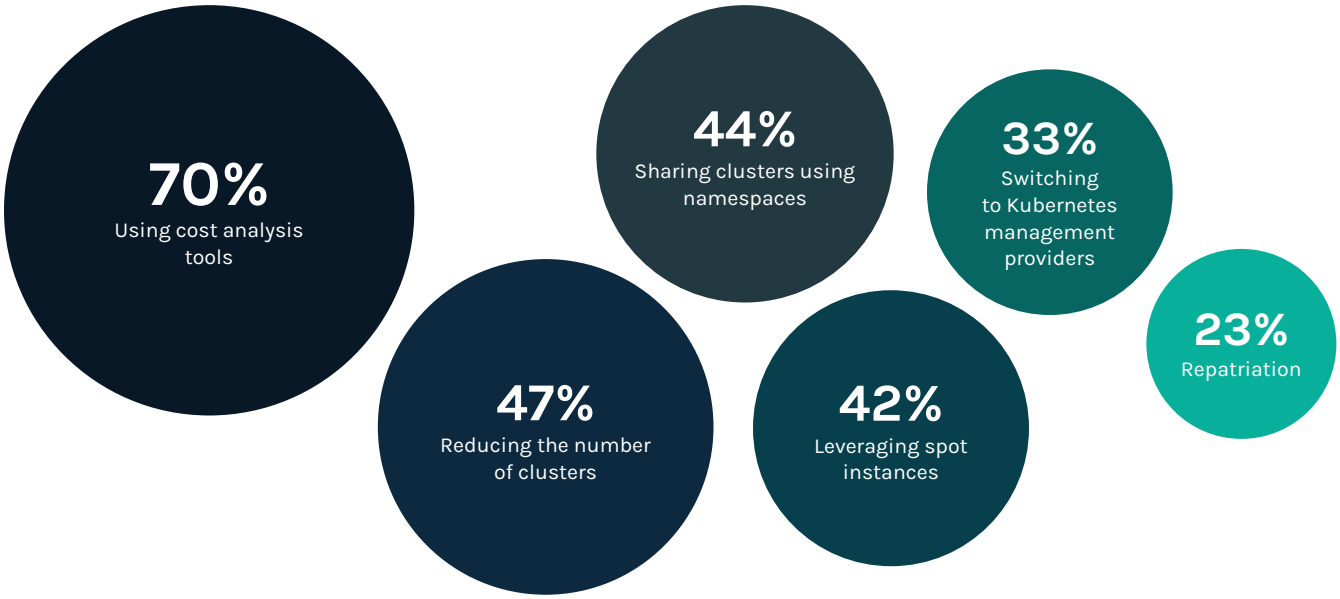

Platform teams are responsible for managing the cost of Kubernetes operations, especially when running multiple clusters spanning various environments. Top cost management priorities for platform teams encompass substantial reduction and optimization of Kubernetes infrastructure and the establishment of transparent visibility and financial accountability in cloud and Kubernetes costs/FinOps.

Current Cost Management Initiatives

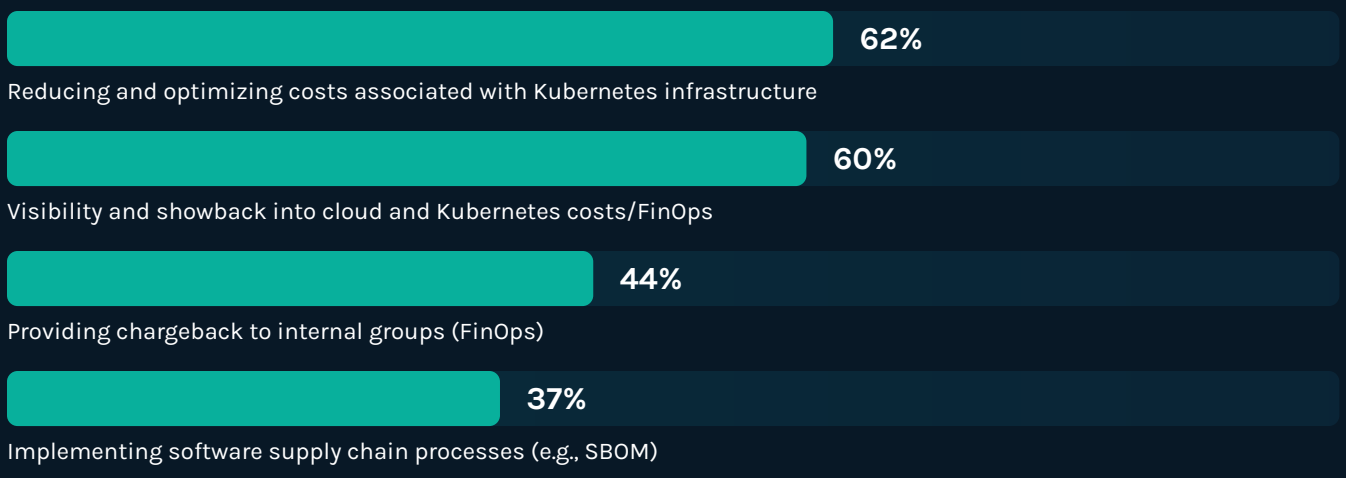

Elevating Infrastructure Management

Platform teams oversee infrastructure management within their organization to ensure standardized tools and workflows are readily available. The rapid pace of digital transformation has prompted enterprises to reevaluate their cloud tech stack and management practices, with many recognizing hybrid and multi-cloud deployments as essential for optimizing their IT infrastructure.

Organizational Preferences in Cloud Infrastructure Models

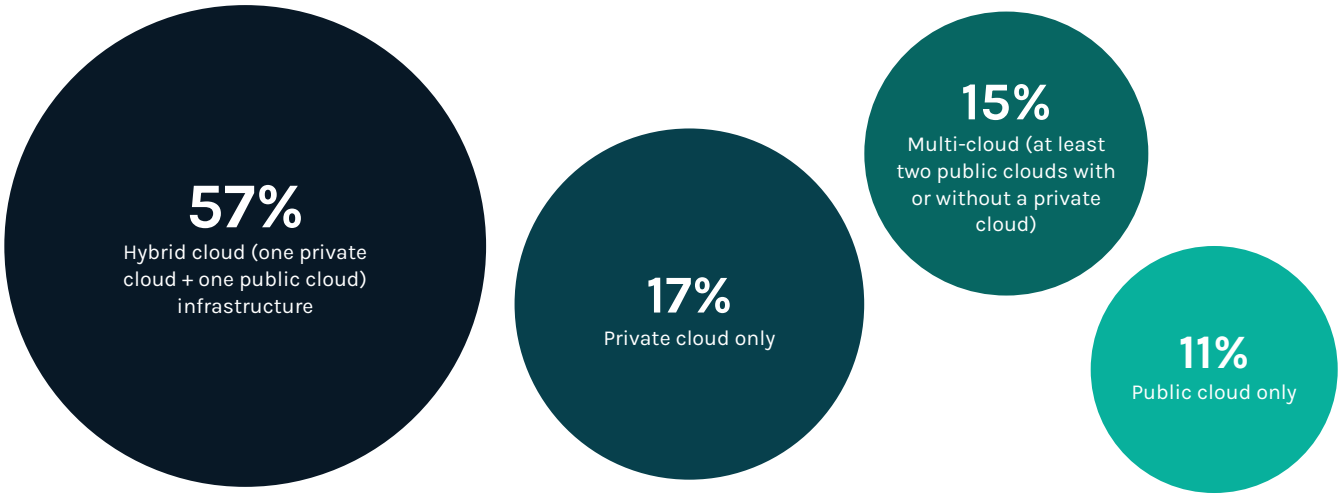

Cloud Infrastructures Used by Organizations

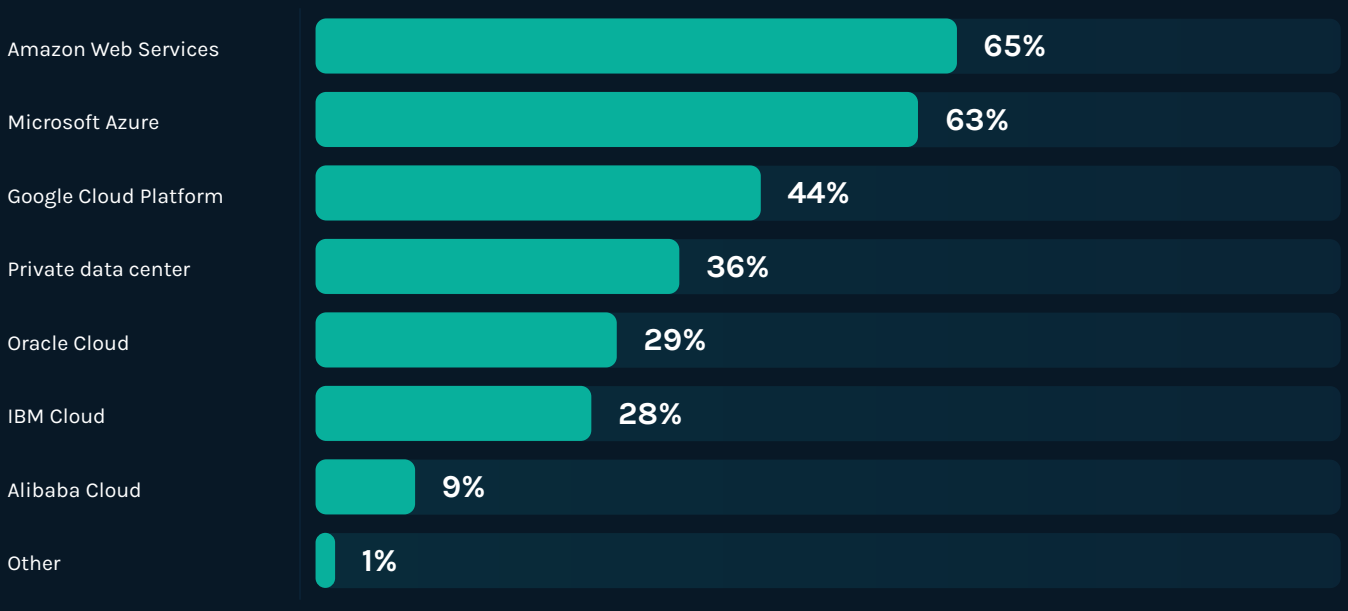

With hybrid and multi-cloud infrastructures more commonly being adopted by enterprises, organizations are prioritizing infrastructure management initiatives including increasing reliability and availability of applications, application and infrastructure observability and monitoring, multi-cloud and hybrid cloud environment provisioning and management, and management of Kubernetes clusters.

Which of the following are current ongoing infrastructure management initiatives for platform teams within your organization?

(Select all that apply)

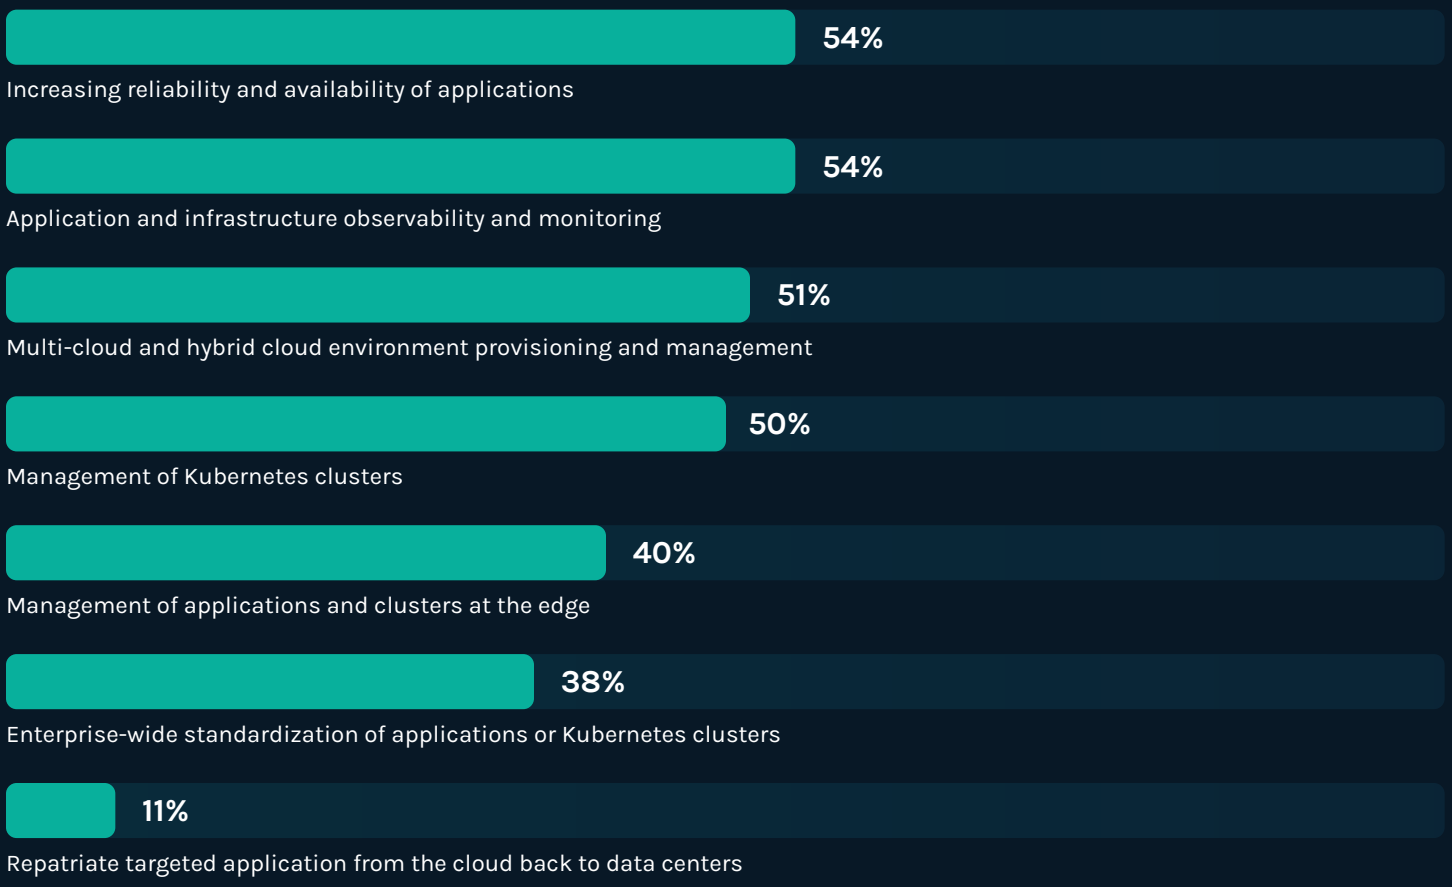

There is a consensus throughout organizations implementing modern applications across the cloud about their struggles to speed innovation while managing high volumes of applications across their cloud infrastructure. Over the past year, the majority of platform teams have experienced pain points as they scale Kubernetes usage with the biggest challenges cited as integrating with existing tools, teams' inclination to build their own tools and managing multi-cloud environments.

Challenges Platform Teams Faced in the Last Year

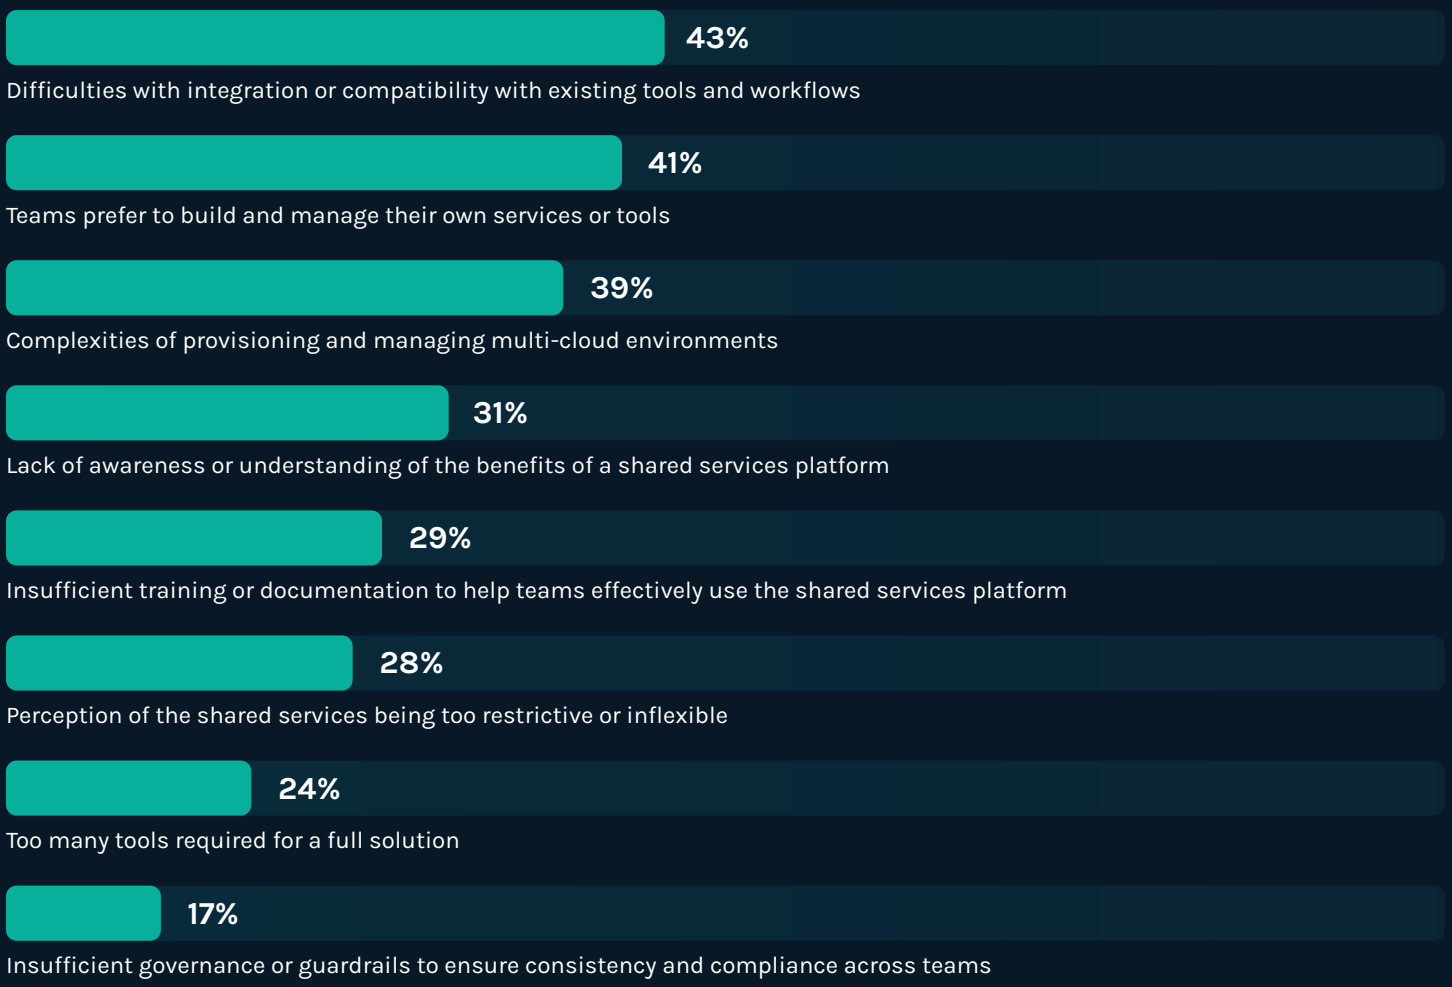

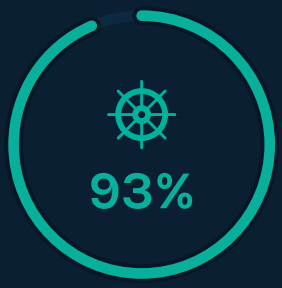

A concerning **majority** of platform teams are currently experiencing challenges, highlighting an opportunity to streamline processes for these teams to reduce the burden of scaling Kubernetes usage.

Driving Developer and Operations Productivity Through Automation and Self-Service

Platform teams are designed to standardize Kubernetes tools and workflows for organizations while supporting and empowering developer and operations teams to move at the pace of innovation. In fact, when asked about developer- and automation-focused initiatives, **almost half** cite the standardization and automation of infrastructure as a priority, in addition to automating cluster provisioning and providing self-service experiences for developers.

Nearly three in four respondents (74%) are focusing on automating cluster provisioning or automating environment provisioning and management.

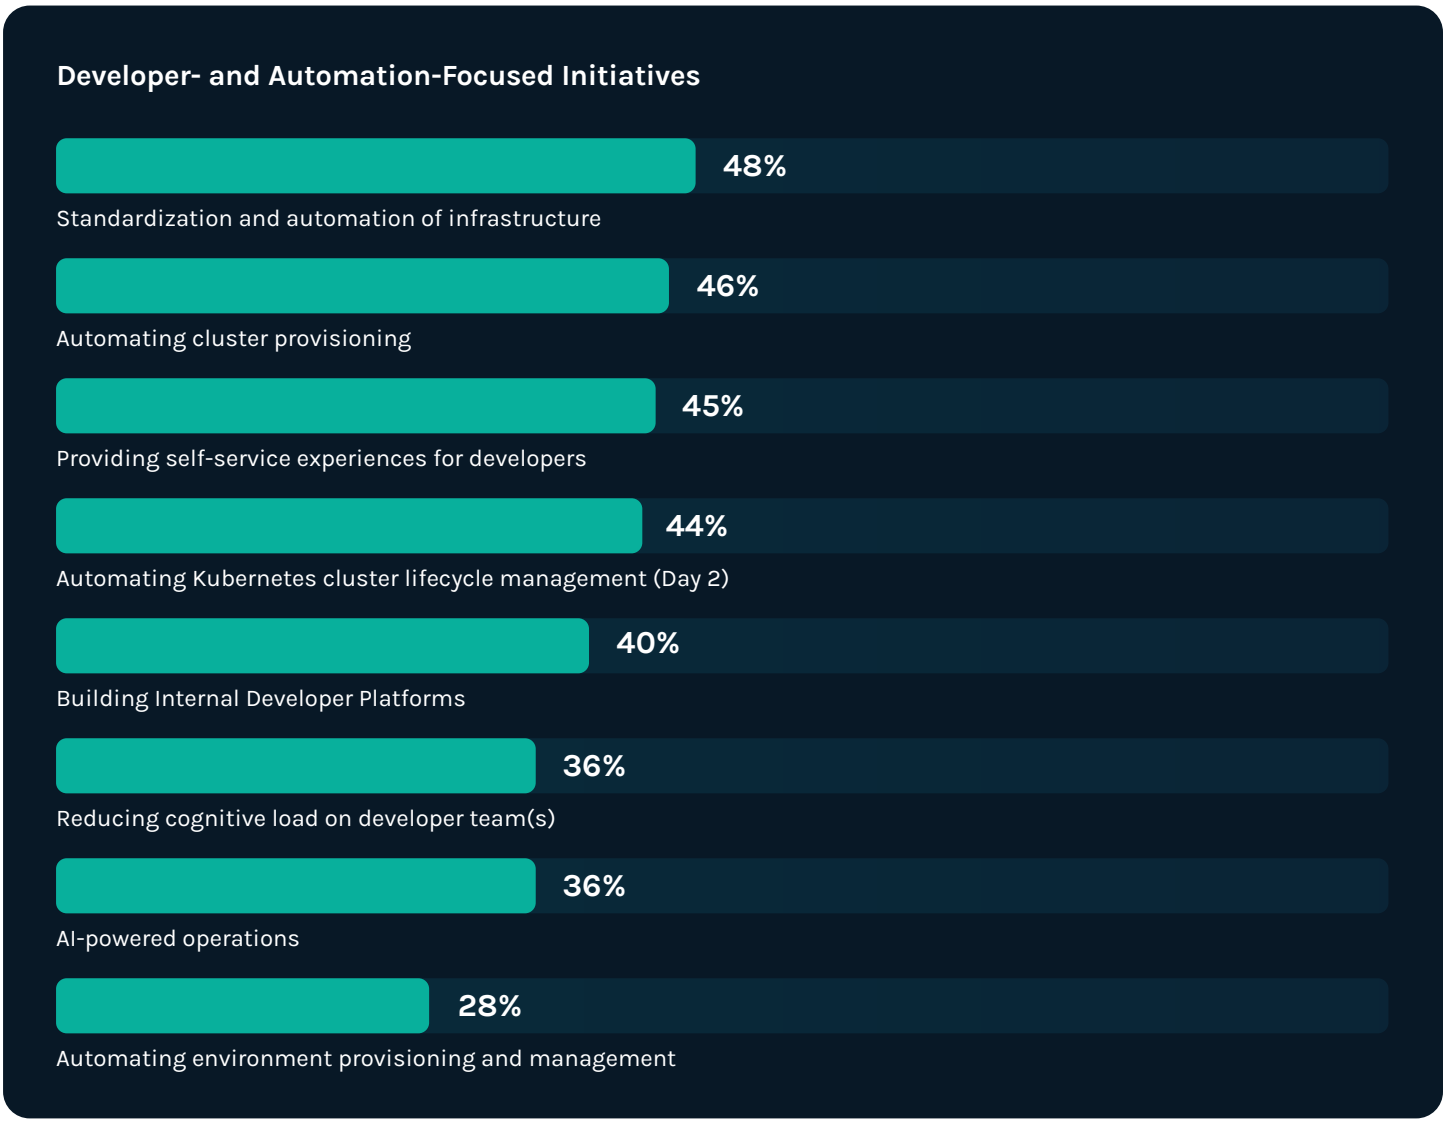

Self-service automation capabilities can accelerate application management for teams by ensuring rapid access to cloud infrastructures. **Of the 45%** who cited providing developer self-service as a top initiative, these organizations have enabled self-service workflows for Kubernetes cluster provisioning, monitoring application and infrastructure performance, namespace provisioning, environment provisioning, CI/CD operations and infrastructure/application updates.

Which of the following Kubernetes-related functions do you currently have enabled as a self-service workflow?

(Select all that apply)

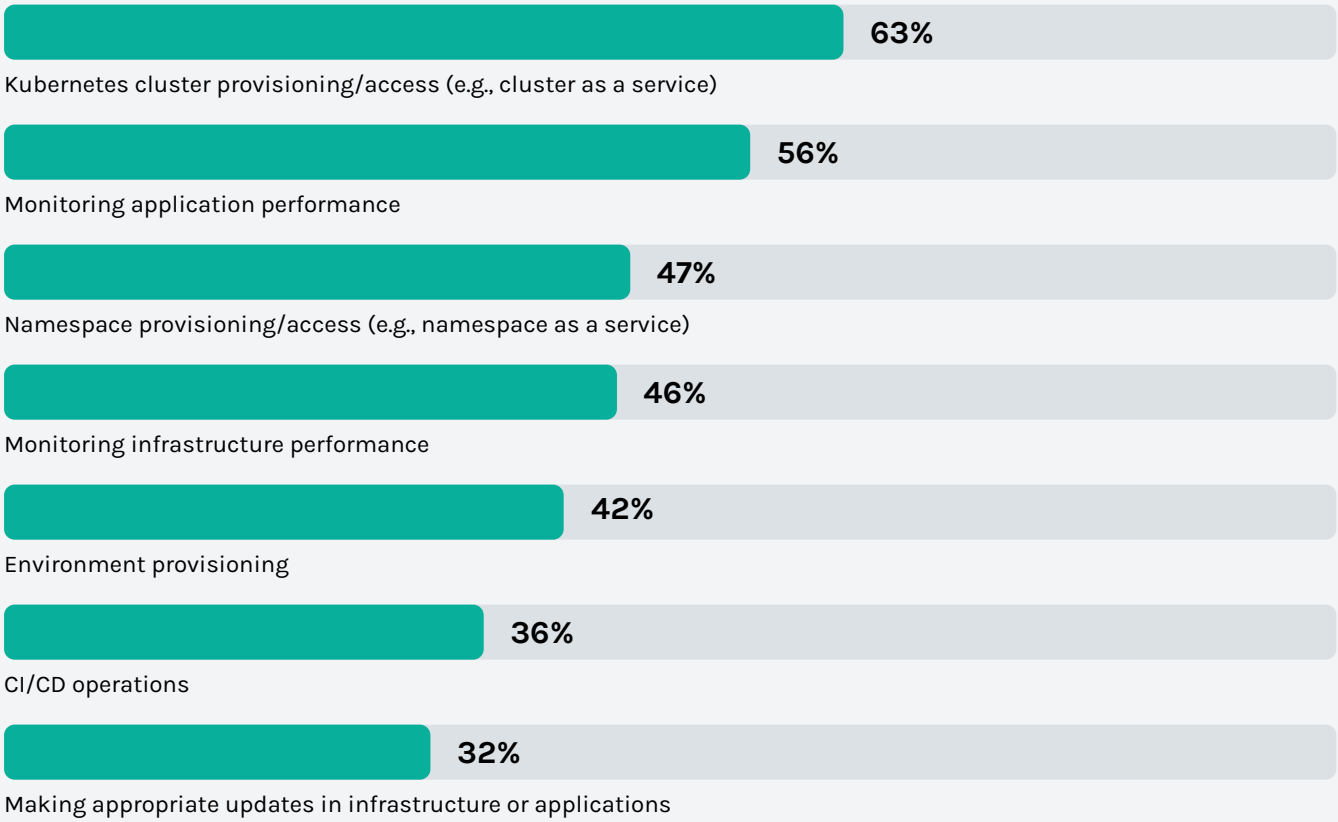

Surge in Kubernetes usage calls for increase in self-service automation

Kubernetes usage has risen globally, with 5.6 million developers using Kubernetes worldwide. As the number of users continues to increase, self-service automation is critical to ensure platform teams and their developers can rapidly deliver cloud applications.

Safeguarding Kubernetes

Platform teams provide company-wide governance guardrails to ensure that Kubernetes clusters and workloads are secure, curbing potential security risks that can increase as Kubernetes usage surges. Top platform team security and compliance priorities include DevSecOps implementation, secure access to clusters and cloud resources, and streamlined compliance and auditing.

Platform Team Security and Compliance Priorities

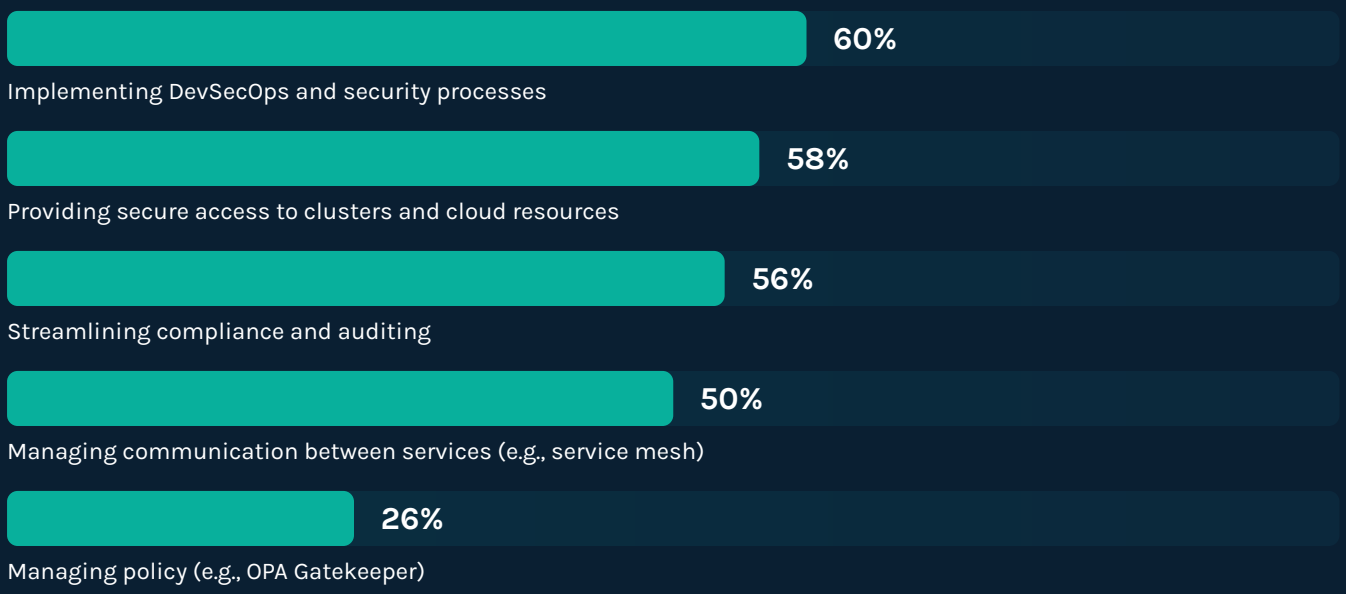

SECTION 3

Future Trends and Challenges in Cloud and Kubernetes Automation

Alongside an exploration of the current trends and challenges organizations face with cloud and Kubernetes automation, the research analyzes projected trends and challenges for the year ahead. It found that organizations are hyper-focused on self-service, AI and automation for developers and platform teams with continued and increased Kubernetes usage anticipated over the next year. With that, organizations plan to prioritize cost allocation, self-service environment provisioning and doubling down on hybrid and multi-cloud for increased efficiency and productivity.

Enterprises Embrace Kubernetes for Empowering Developers and Elevating Tech Stacks

Almost all **(95%)** organizations plan to increase Kubernetes usage in the next year. Specifically, seven in 10 organizations **(71%)** plan to increase Kubernetes usage anywhere from **26% to 100%**.

In the next year, how much does your organization plan to increase its usage (approximately by number of clusters)?

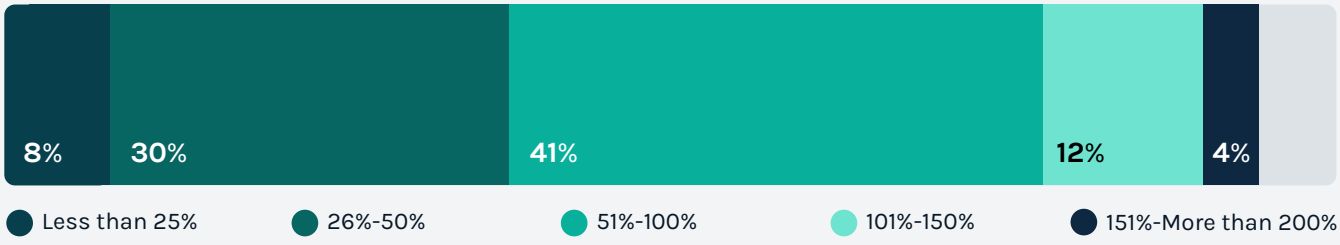

With continued increased Kubernetes usage, organizations are gearing up to empower their platform teams to provide developers with the tools they need.

Boosting Developer and Platform Team Productivity with Self-Service, AI and Automation

Platform engineering teams aim to help developers move more quickly and efficiently with the right tools and guardrails in place. Automation, standardization and self-service experiences are several methods platform teams can leverage to accelerate developer productivity and alleviate tensions between operations and platform teams.

Organizations plan to prioritize the following developer- and automation-focused initiatives in the next year:

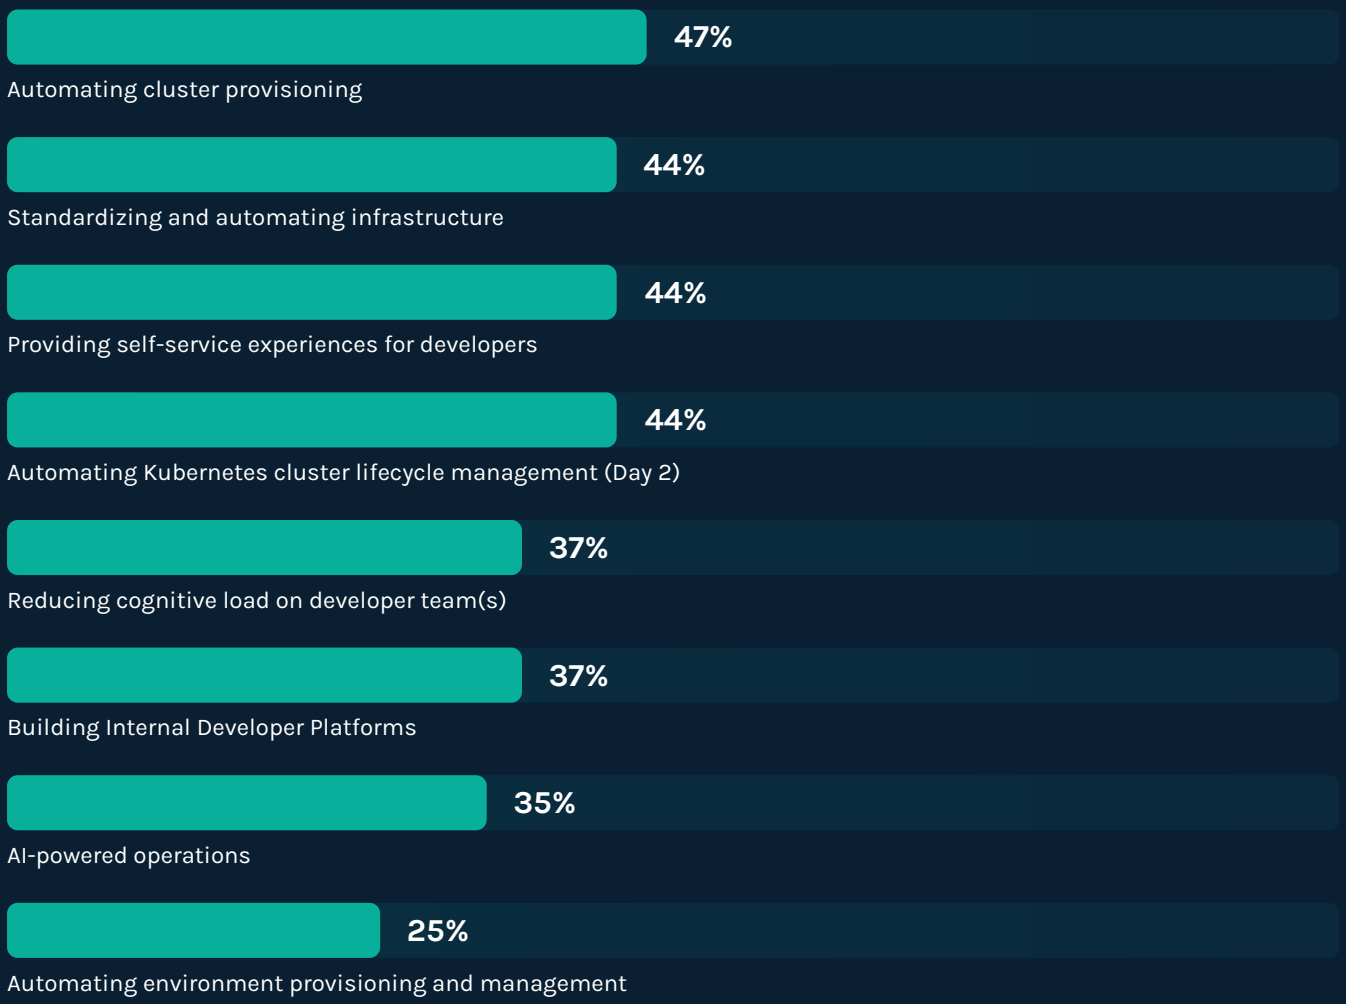

Of the 44% who plan to prioritize developer self-service, the top functions platform teams would like to enable as self-service workflows include:

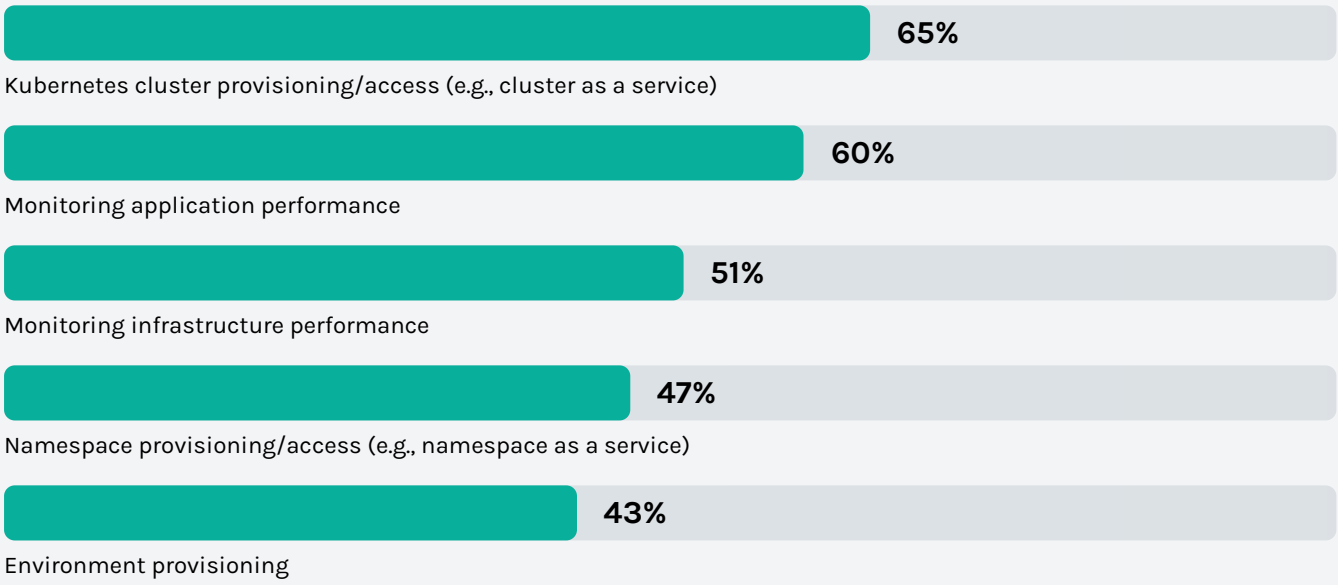

Revolutionizing Environment Provisioning for Improved Developer Experience

Environment provisioning is mired in manual processes and infrastructure dependencies that affect the developer experience. When asked about challenges platform teams will face in the next year, **39%** believe the complexities of provisioning and managing multi-cloud environments are a top inhibitor.

There's a clear gap between the perceived importance of environment provisioning and the satisfaction with existing processes. The dissatisfaction could be indicative of a growing recognition that traditional approaches may not be sufficient, especially in the context of emerging challenges like managing multi-cloud environments.

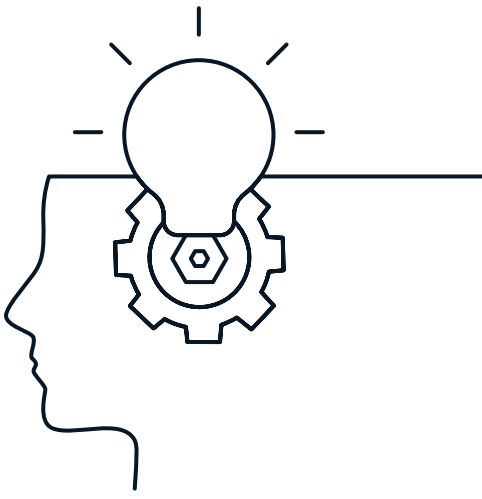

For example, developers face challenges with environment provisioning that hinder their ability to deploy an application or service from code-complete to production in their ideal amount of time. An automated environment provisioning solution with a self-service capability can abstract the complexities and decrease the time it takes to provision and access Kubernetes-based environments.

28% Over a quarter cited automating environment provisioning and management as a current developer and automation-focused initiative for platform teams.

A vast majority of respondents who automate environment provisioning and management highlighted the extensive time it currently takes to deploy a cloud application/service from code-complete and quality assurance to production — anywhere from one month to longer than one year.

Current time for deploying a cloud/ application service

65%
one to three months

17%
six months to longer than a year

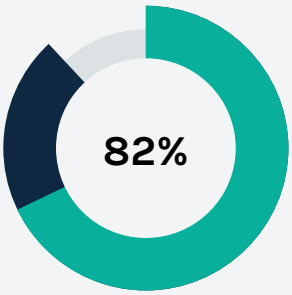

What is causing the delay from code-complete to production?

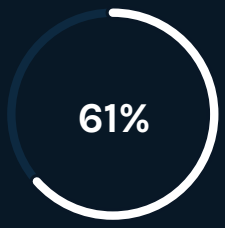

There are endless cycles of troubleshooting at the infrastructure level prior to deployment

Environment provisioning and management is too complex

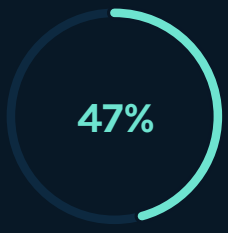

There is too much back and forth between the application and infrastructure teams to determine the optimal infrastructure required for the applications

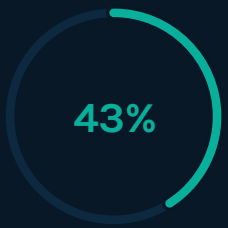

They have to wait on someone else to provision resources or spin up environments

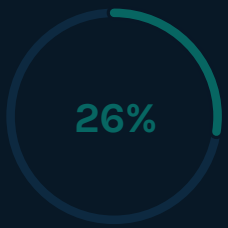

Ticking systems for provisioning are too onerous

According to **80%** of these same respondents, deploying a cloud application/service should take within hours up to six days.

Ideal time for deploying a cloud/application service

Provisioning and managing multi-cloud environments are top inhibitors for platform teams, making it critical for organizations to determine how best to work within various cloud infrastructures.

Of the 25% of respondents advocating for automated environment provisioning and management as a future developer and automation-focused initiative:

prioritize empowering developers to independently deploy and manage environments.

see value in a self-service workflow/portal for developers to provision environments autonomously.

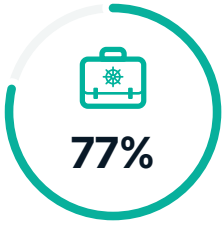

prioritize endorsing developers' ability to deploy and manage Kubernetes clusters themselves.

champion the implementation of a self-service workflow/portal for developers to provision Kubernetes clusters themselves.

Organizations are Doubling Down on Hybrid and Multi-cloud

Hybrid cloud will continue to be the preferred mode of deployment for nearly half of organizations as they seek a balance between on-premises and cloud-based deployments to meet their evolving needs — **49%** of organizations plan to use hybrid cloud (one private cloud + one public cloud) in the next year. This trend underscores the growing recognition that a thoughtful mix of on-premises and cloud-based solutions is essential for meeting the evolving needs of modern businesses in an increasingly digital landscape.

Organizations are actively exploring and adapting to the dynamic possibilities offered by multi-cloud architectures. There is growing interest in multi-cloud suggesting a notable shift in strategy — more than a quarter **(28%)** plan to leverage a multi-cloud infrastructure (at least two public clouds with or without a private cloud) in the next year, compared to **15%** who currently use it. This could signify a strategic response to the complexities of the modern digital landscape and the growth of generative AI, where diversification of cloud resources is seen as a proactive and forward-thinking approach.

In the next year, AWS, Microsoft Azure and Google Cloud Platform will continue to be the top infrastructures used by organizations.

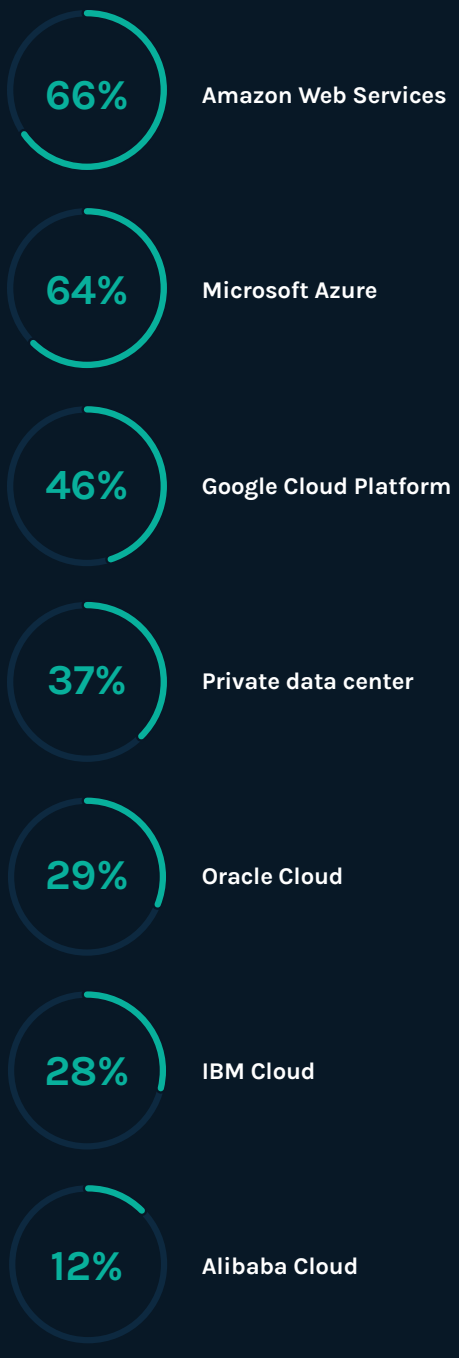

However, when working among various cloud infrastructures, it can be incredibly difficult to view an organization's cloud spend across infrastructures and report how spend is allocated across internal departments. With many enterprises overspending in the cloud and a renewed focus on business efficiency, there is an opportunity for organizations to define cost management initiatives to accurately view Kubernetes cloud spend across clusters.

In the next year, platform teams will prioritize FinOps and reduce Kubernetes infrastructure costs.

Planned Cost Management Initiatives for the Next Year

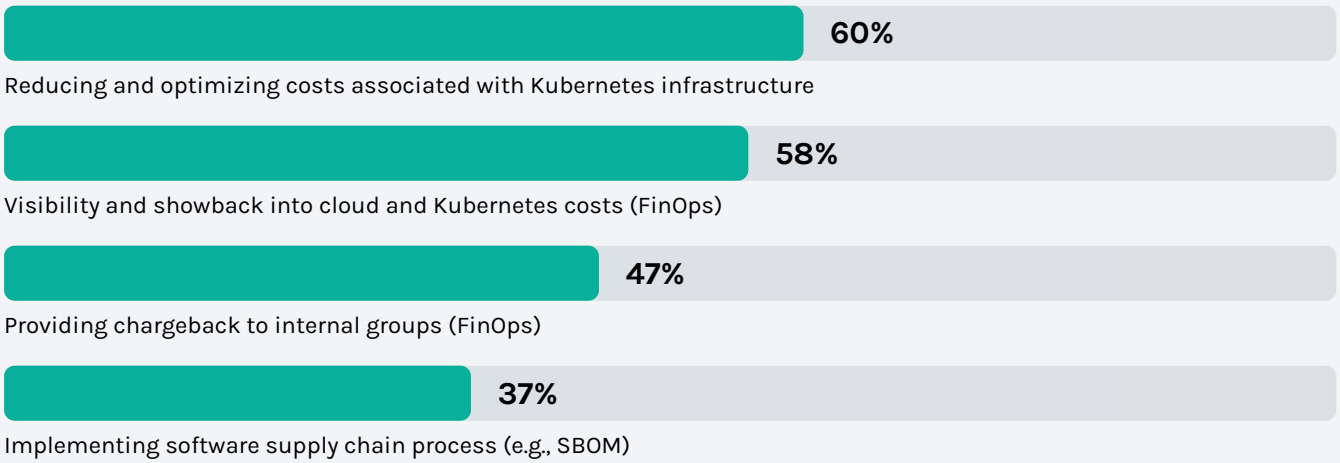

Standardizing Hybrid and Multi-cloud Environments

The escalating complexity of standardizing and managing applications across hybrid or multi-cloud environments — coupled with the limitations of existing automation tools — has left platform, developer and operations teams with breach-prone and time-consuming approaches. Role-based access control is one solution, however, having to configure and manage it cluster by cluster is a burden. Platform teams need to streamline and consolidate access control for clusters spanning different operating environments, clouds and on-premises data centers.

In the next year, platform teams will prioritize simplifying security and compliance processes.

Which of the following are future security/compliance initiatives planned in the next year for platform teams within your organization?

(Select all that apply)

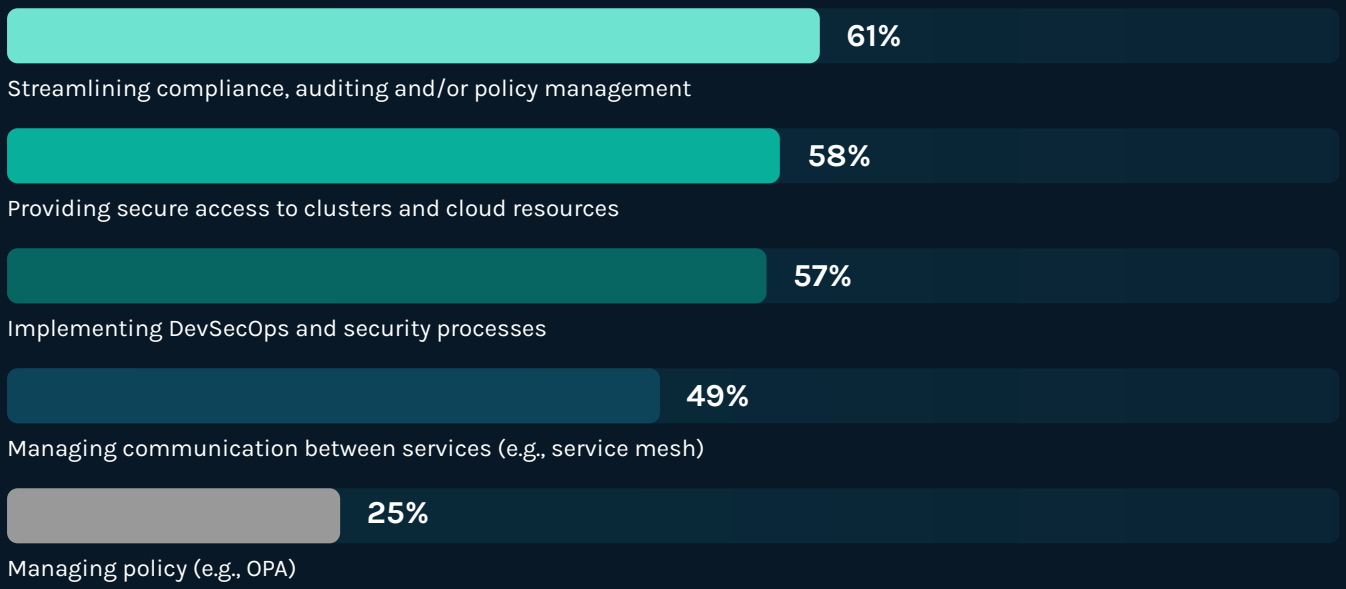

SECTION 4

Emerging Frontier: AI and GenAI Initiatives in the Enterprise

The Imperative of AI and GenAI

As organizations continue to embrace cloud technologies and Kubernetes, they are simultaneously investing heavily in AI and GenAI capabilities. As with cloud and Kubernetes adoption, organizations recognize the potential of these technologies but face significant challenges in implementation and management. Platform teams, having proved their value in navigating the complexities of cloud and Kubernetes, are now poised to play a crucial role in AI and GenAI adoption.

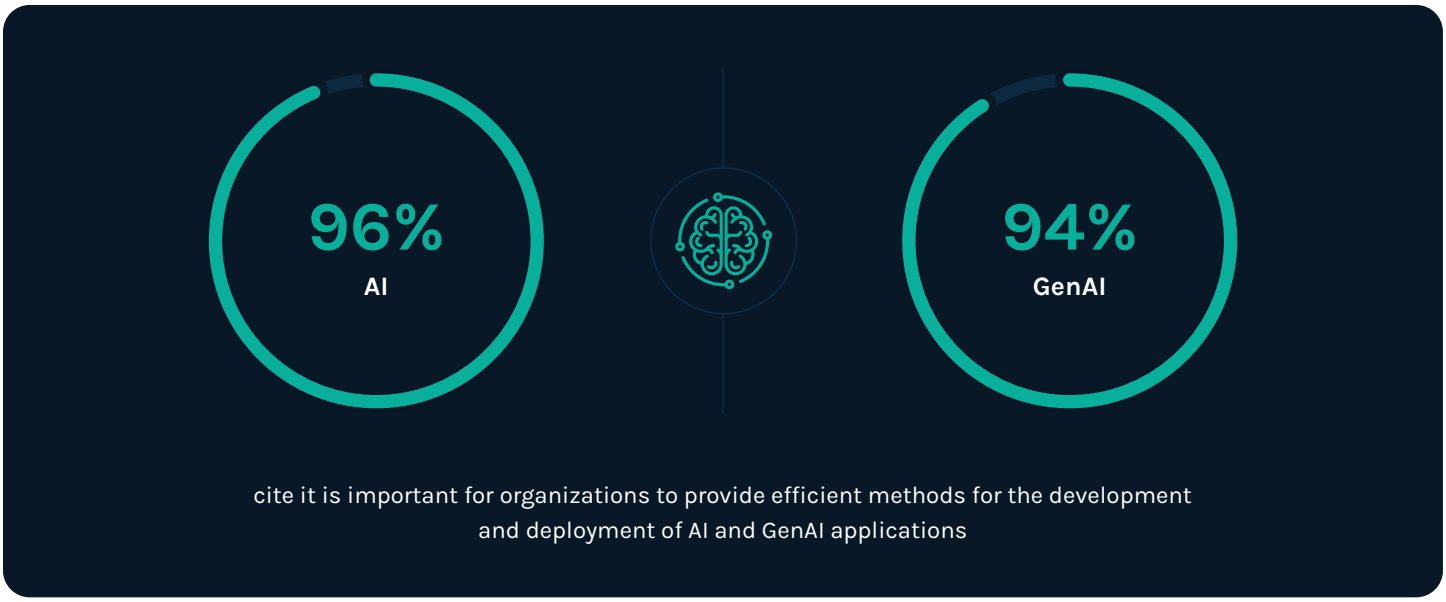

The near-universal recognition of AI and GenAI's importance underscores a significant shift in enterprise technology strategies. Organizations are no longer viewing AI as a future possibility but as a current necessity. This aligns closely with the broader trends in cloud and Kubernetes adoption, where companies are seeking to modernize their infrastructure to support more advanced and demanding technologies.

Current State of AI Operations

Organizations are at various stages in their AI journey, particularly in the implementation of Machine Learning Operations (MLOps) and Large Language Model Operations (LLMOps).

Of those that found it important for their organization to provide efficient methods for the development and deployment of AI applications, or plan to revisit it in the next 12 months, **94%** have implemented or adopted MLOps. Additionally, **nine out of 10** companies that emphasize the importance of efficient GenAI development processes have already implemented LLMOps practices, leveraging either public or private large language models.

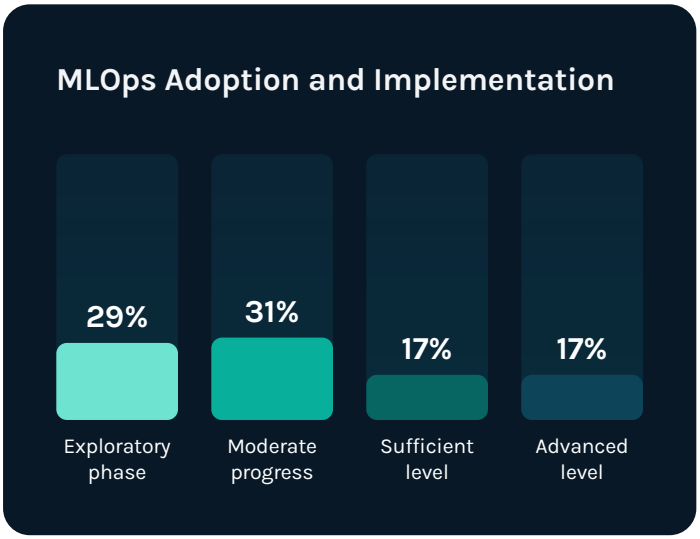

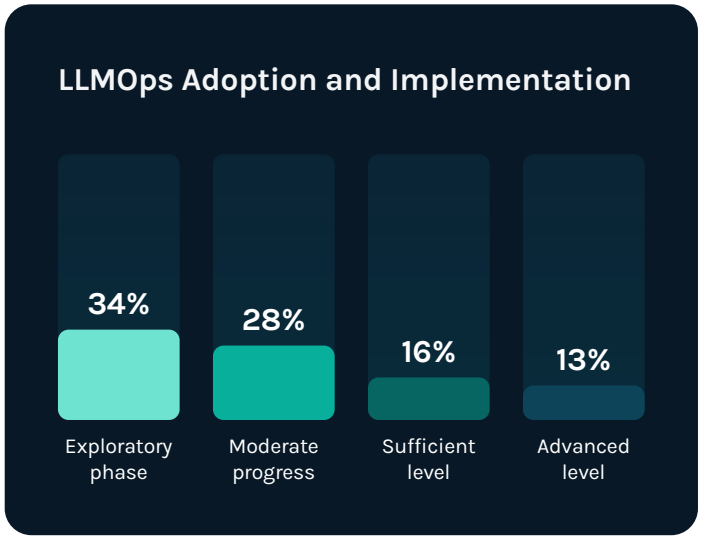

While many organizations have begun their AI operations journey, the majority are still in the early to middle stages of adoption. Specifically, **less than a quarter** of organizations cite being at a sufficient level of adoption and implementation of both MLOps and LLMOps. This mirrors the Kubernetes adoption curve, suggesting that AI operations are following a similar pattern of gradual integration into enterprise environments, potentially due to roadblocks that make the implementation challenging.

Challenges in AI Implementation

Despite the recognized importance of AI and GenAI, organizations face significant hurdles in implementation. These challenges echo many of the difficulties encountered in cloud and Kubernetes adoption, particularly around integration, complexity and management.

Of those who have implemented MLOps on some level, **nearly all (95%)** cite it is currently challenging for their developers to experiment with and deploy AI apps. Similar findings were reported for experimenting with and deploying GenAI apps — **94%** cite it is currently challenging for their developers.

Top MLOps Challenges:

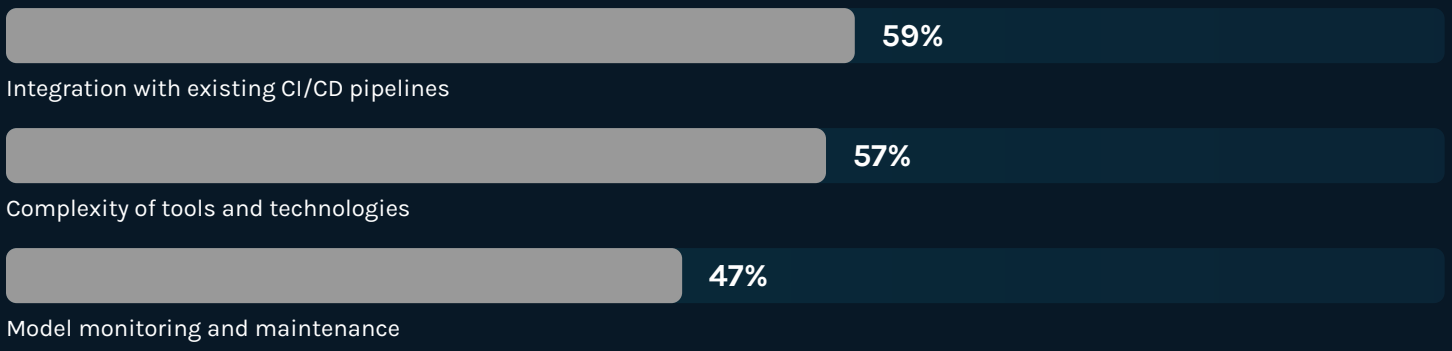

Top Public LLMs Challenges:

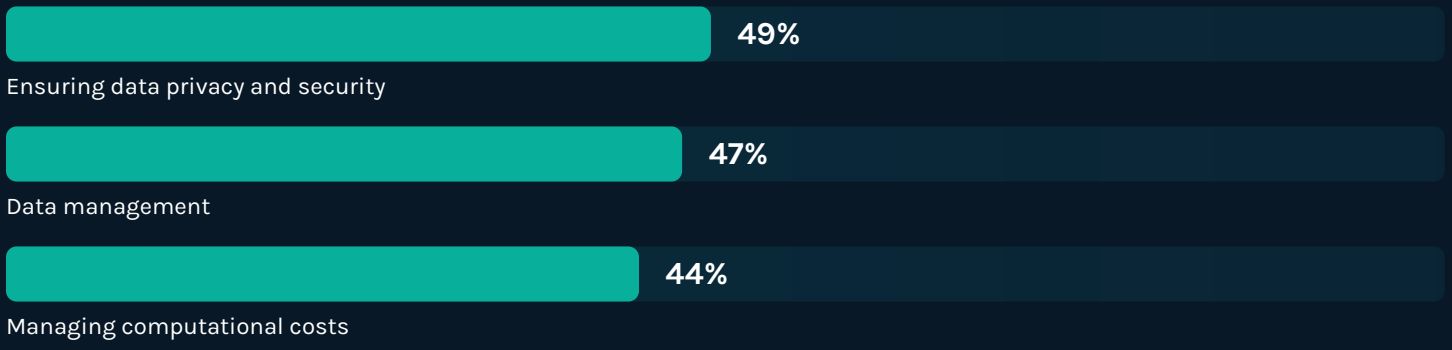

Top Private LLMs Challenges:

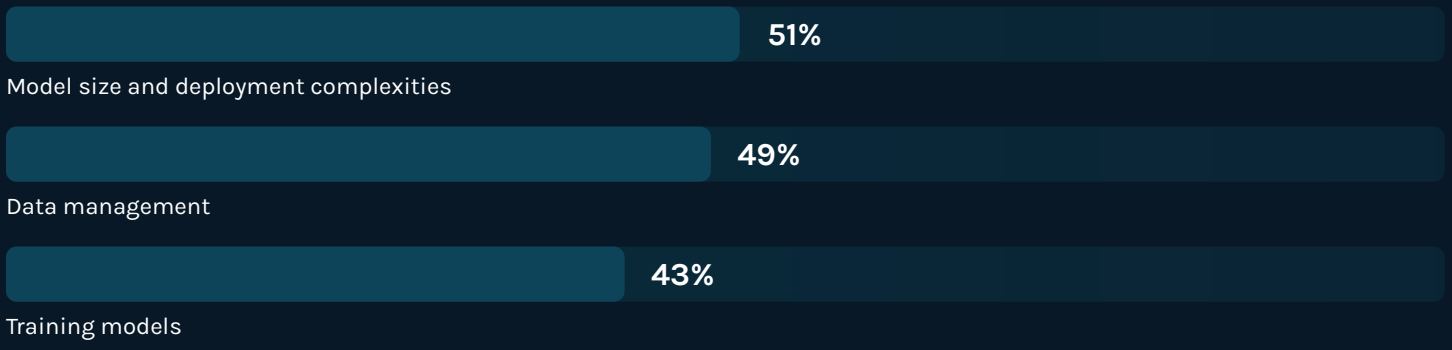

The prevalence of integration and complexity challenges underscores the need for streamlined, end-to-end solutions that can seamlessly incorporate AI and GenAI capabilities into existing development and operational workflows. This parallels the broader industry trajectory towards automated, self-service solutions in cloud and Kubernetes landscapes.

The Role of Platform Teams in AI Initiatives

As organizations grapple with these AI challenges, platform teams are emerging as key enablers of AI adoption and management, much as they have for cloud and Kubernetes technologies.

Top responsibilities respondents believe should be assigned to platform teams to assist in the development of AI/GenAI applications:

Security for MLOps and LLMOps workflows

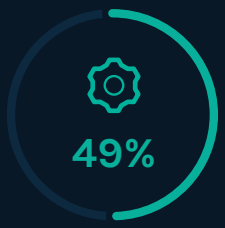

Model deployment automation

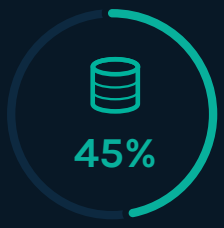

Data pipeline management

The focus on security, automation and data management highlights the critical role platform teams play in creating a stable, efficient and compliant foundation for AI initiatives.

Enhancing Developer Productivity Through AI-Focused Self-Service

The importance of developer productivity and self-service capabilities extends to AI initiatives as well — **83%** of respondents believe that self-service access to pre-configured AI workspaces with all the necessary MLOps and LLMOps tooling could save their developers and data scientists **10%** or more time each month.

Estimated developer time savings with AI-focused self-service:

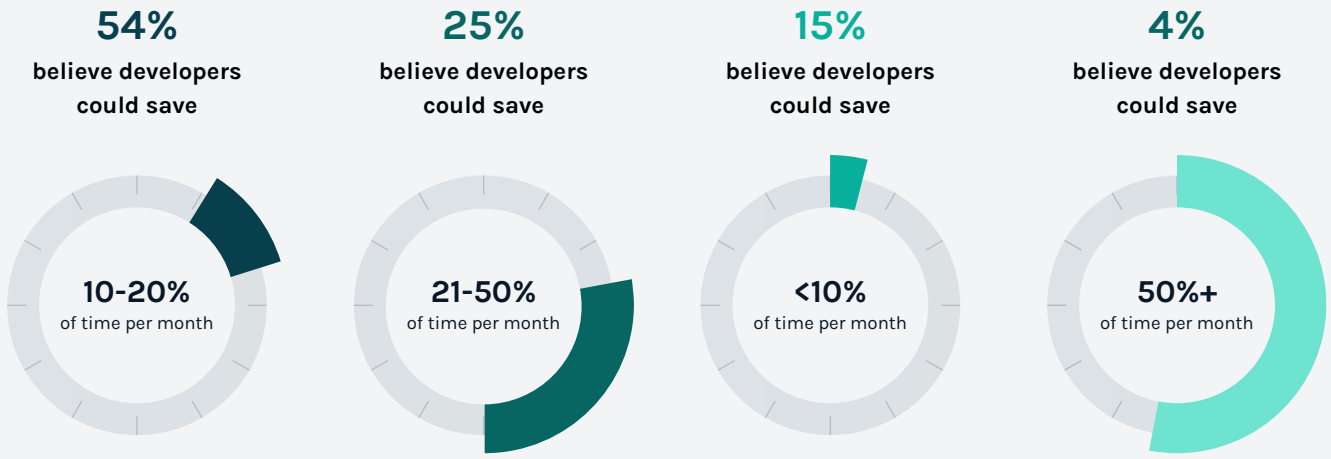

These **potential time savings are significant**, especially when considered in the context of the broader push for developer efficiency we've seen in cloud and Kubernetes management. By providing self-service access to pre-configured AI workspaces, platform teams can dramatically reduce the cognitive load on developers, allowing them to focus on innovation rather than infrastructure management. Organizations could see an overall increased output from developers and data scientists, yielding higher return on investment.

For example, in an organization with 100 developers earning an average salary of $140,000, adopting self-service AI workspaces could unlock nearly 20,000 hours of developer time annually.

This is equivalent to **$1.4 million** in salary costs or the productivity gain of 9 additional full-time developers — without increasing headcount.

The Future of AI in Enterprise Environments

Looking ahead, organizations are planning to invest in solutions that streamline AI and GenAI management, much as they have done with cloud and Kubernetes technologies.

Key capabilities valued by organizations include:

- Pre-configured environments for developing and testing generative AI applications with LLMs
- Automatic allocation of AI workloads to appropriate GPU resources based on cost and performance
- Pre-built MLOps pipelines that standardize and streamline the model development lifecycle
- Ability to virtualize and share GPUs across multiple AI workloads to improve utilization and lower costs
- GPU matchmaking to dynamically match the user workspaces with available GPUs or pools of GPUs based on criteria such as proximity, cost efficiency, GPU type and more to improve utilization

The findings from this AI-focused survey align closely with the broader trends we've observed in cloud and Kubernetes adoption, highlighting the interconnected nature of these cutting-edge technologies in modern enterprise environments. These insights offer a comprehensive view of the evolving technology landscape that platform teams must navigate.

SECTION 5

Current and Future Trends of Platform Teams

As organizations increasingly adopt cloud, Kubernetes, AI and GenAI technologies, platform engineering teams play a vital role in accelerating innovation, improving operational efficiency and ensuring consistency throughout the software development lifecycle.

Platform Teams Can Boost Organizational Productivity with the Right Tools and Resources

Organizations see the value in platform teams — **nine in 10 (91%)** have a defined platform engineering team or a group acting as a platform team. Out of those that do not currently have a platform team, **87%** plan to build a dedicated team in the next year.

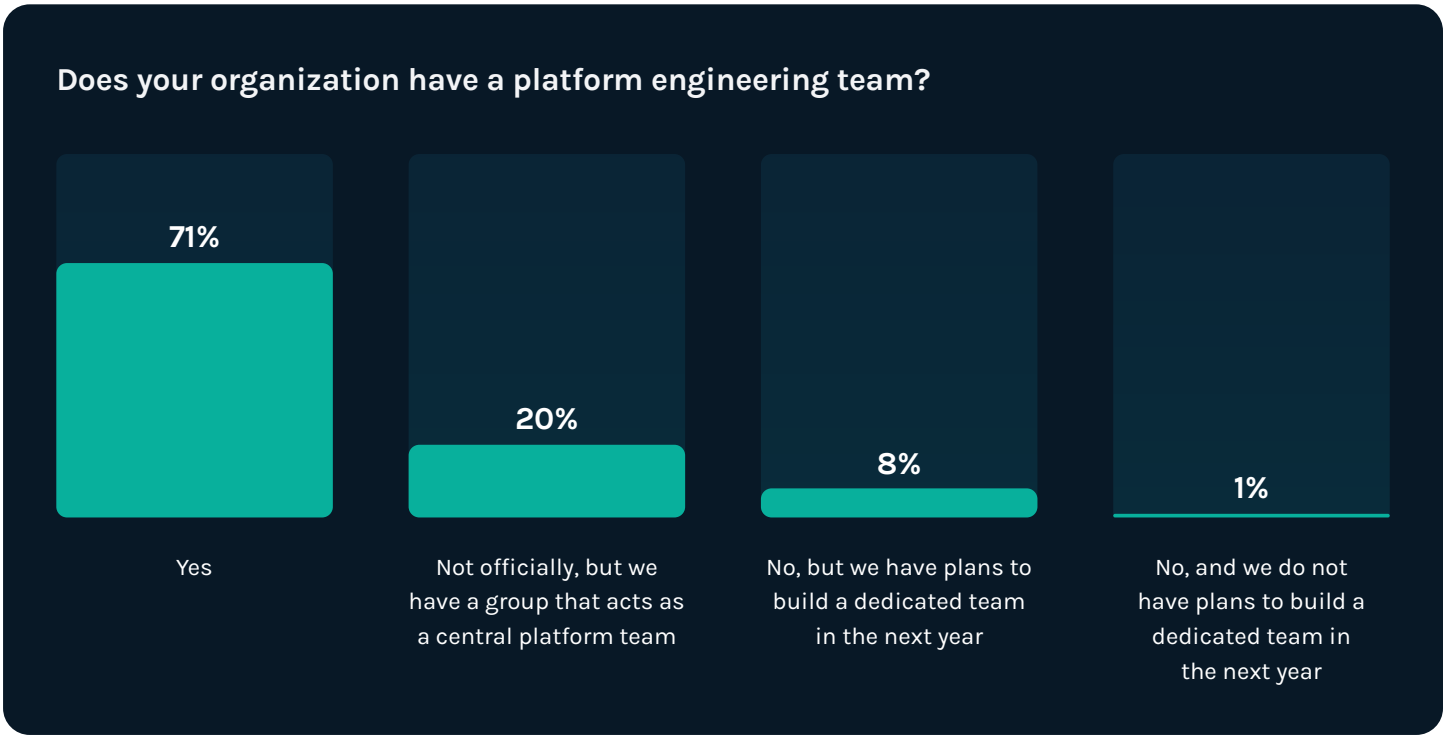

There is a diverse landscape in terms of organizational resources and priorities, with possible implications for how teams approach and execute platform engineering initiatives. A significant portion of organizations have a moderately sized platform team — **nearly half of respondents (49%)** said their platform engineering team consists of six to 15 people while **nearly a third (32%)** have a larger team of more than 15 people. This could suggest that these teams are staffed to handle complex tasks and projects.

On the other hand, **less than two in 10 (18%)** organizations have a small platform engineering team of one to five people, opening up the possibility that these teams could face some challenges in terms of bandwidth and resources. **Almost everyone within this group (86%)** wants to build out a dedicated platform team in the next year.

Roles represented within platform teams:

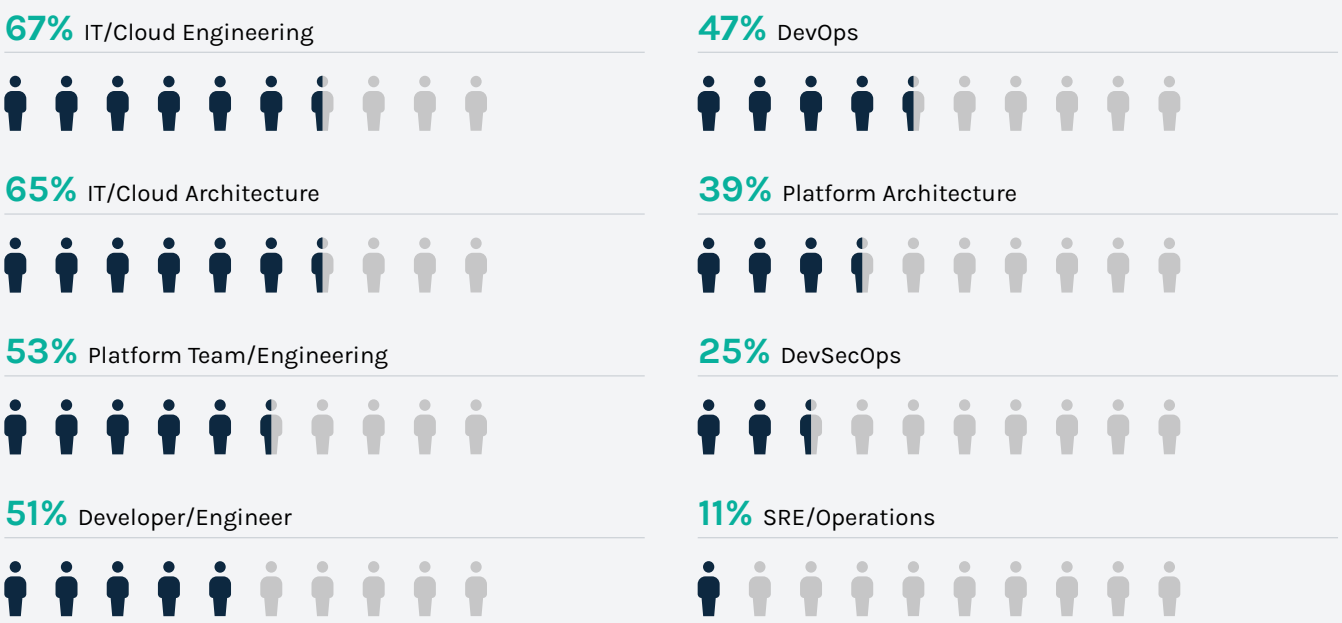

These differentiating roles within platform teams contribute to the same projects. However, working on the same project usually means sharing the same resources and budget to be split between a diverse set of needs across cloud infrastructure, Kubernetes and security. Individual engineering teams can overprovision resources without knowing the impact to other teams and the company's overall budget. Organizations often find that Kubernetes resource consumption and spend can rapidly increase — soaring past spend limits.

While some platform teams share a budget with IT or don't have a budget at all, other organizations give their platform teams financial autonomy by providing separate budgets, which could indicate a strategic move, allowing these teams more control over their resources and decision-making process.

Platform Team Budget Allocation

53%
share a budget with IT or do not have their own budget at all

47%
have a platform team with its own budget separate from IT

Of organizations with a platform team budget, **more than half** of budgets range from $500,000 to upwards of $1.9 million, suggesting that organizations are willing to invest significantly in their platform teams.

What is the size of your platform team's engineering budget?

Platform teams, with ample budgets, can allocate resources to drive initiatives including developer and automation-focused efforts, cost management, security and compliance, and infrastructure management. There is a recognition of the importance of platform teams, not just in giving them autonomy but also in allocating substantial financial resources. This could imply a heightened focus on technological infrastructure and innovation within these organizations.

Platform Teams are Critical Enablers of Innovation and Efficiency for Developers

For successful innovation to occur, platform teams and developers must work together harmoniously — and doing so requires a set of familiar, reusable solutions and capabilities that provide self-service, automation and AI, as outlined by the findings. Yet, platform teams are stretched to their limits managing complex multi-cluster Kubernetes environments. Spending excessive time on manual Kubernetes operations diverts their attention from tasks that could otherwise propel and expedite strategic business objectives.

Nearly all respondents (93%) expect platform teams to face a number of roadblocks in the future, many of which are still the same as they are currently experiencing.

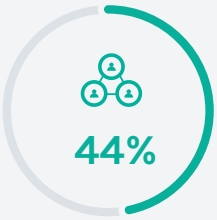

Perception of the shared services being too restrictive or inflexible

Difficulties with integration or compatibility with existing tools and workflows

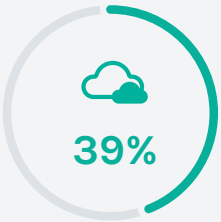

Complexities of provisioning and managing multi-cloud environments

Other challenges organizations believe platform teams will face in the next year:

Teams prefer to build and manage their own services or tools

Lack of awareness or understanding of the benefits of a shared services platform

Insufficient training or documentation to help teams effectively use the shared services platform

Too many tools required for a full solution

The rapid pace of needed changes

Insufficient governance or guardrails to ensure consistency and compliance across teams

Developer and operations teams heavily rely on platform teams to keep processes streamlined and relieve them of the burden of Kubernetes complexities.

When asked how organizations expect the size of their platform engineering team to change over the next year, 87% expect their team to grow:

28% expect **35-59%** growth

28% expect **25-34%** growth

18% expect **0-24%** growth

13% expect more than **60%** growth

8% said their current size is appropriate

When asked how organizations expect their platform engineering budget to change over the next year, 88% expect their budget to grow:

23% expect **0-24%** growth

30% expect **25-34%** growth

24% expect **35-59%** growth

11% expect more than **60%** growth

11% said their current size is appropriate

In order to take the pressure off of platform teams, organizations must provide proper tools and resources to enable them to efficiently make decisions, optimize resource utilization and ensure consistent deployment practices across the organization.

To boost productivity when working in infrastructures, platform teams are gearing up to tackle a range of critical infrastructure management initiatives, with an unwavering focus on enhancing application reliability, bolstering observability and monitoring, and ensuring seamless provisioning and management of hybrid and multi-cloud environments.

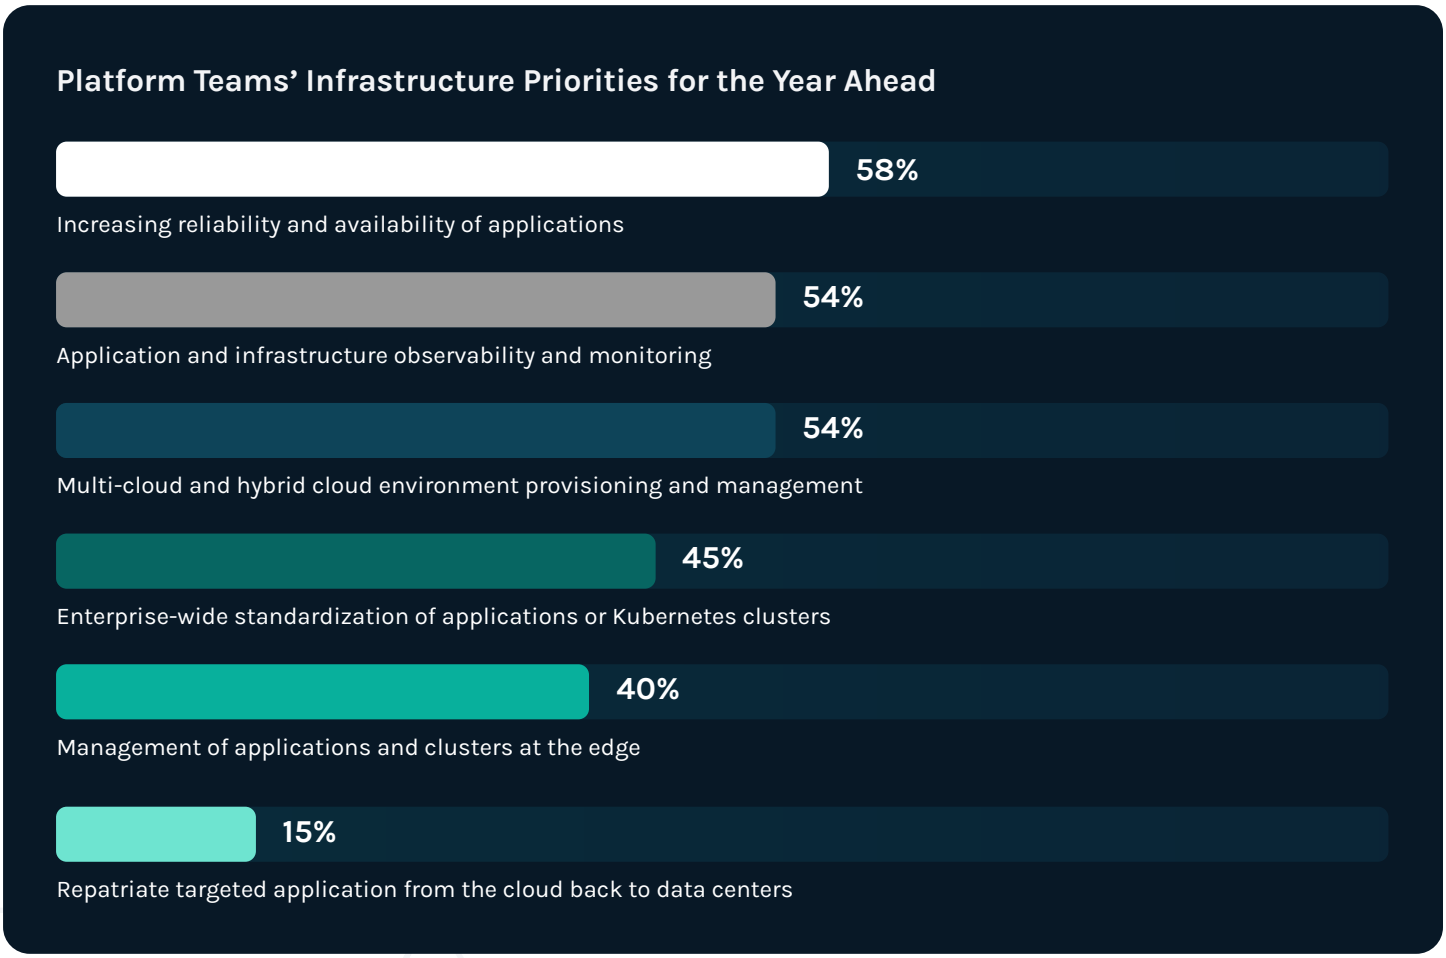

Conclusion

The research underscores the critical role of platform teams in modern enterprise technology environments. As organizations grapple with the complexities of cloud-native architectures, Kubernetes and emerging AI technologies, platform teams are becoming central to successful implementation and management. These teams are uniquely positioned to address cross-cutting concerns such as security, scalability and cost optimization across diverse technology stacks.

The challenges faced by platform teams — from tool integration to multi-cloud management — highlight the need for robust, versatile tooling. Self-service capabilities emerge as a key solution, offering the potential to significantly boost developer productivity while maintaining organizational control. By providing streamlined access to both Kubernetes resources and AI workspaces, platform teams can reduce operational friction and accelerate innovation cycles.

The data suggests that organizations investing in and empowering their platform teams are better equipped to navigate the rapidly evolving technology landscape. As the boundaries between traditional infrastructure, containerization and AI continue to blur, the adaptability and expertise of platform teams become increasingly valuable. Their ability to implement standardized, efficient processes across these domains will likely be a determining factor in an organization's technological agility and competitiveness in the coming years.

Survey Respondent Demographics

1,050 U.S.-based respondents were surveyed to explore current and future trends and challenges associated with Kubernetes and Kubernetes automation.

Primary Role

12% Platform Team/Engineering/Operations

3% Platform Architecture

25% IT/Cloud Architecture

27% IT/Cloud Engineering

20% Developer/Engineer

9% DevOps

4% SRE/Operations

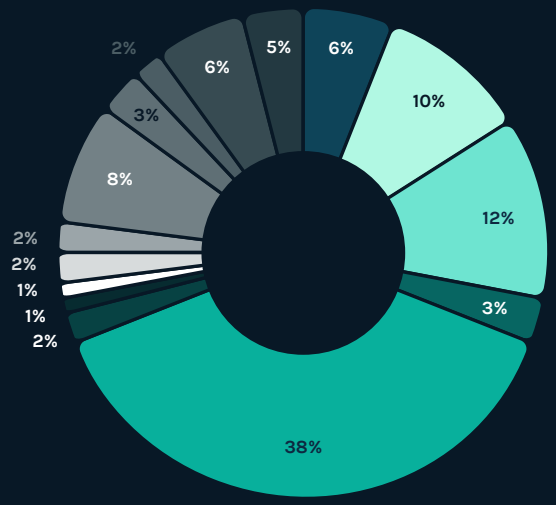

Industry

- Healthcare
- Financial Services
- Manufacturing
- Education
- Computer Software/Hardware
- Media & Entertainment
- Marketing & Advertising
- Hospitality
- Public Administration/Government
- Energy
- Retail & CPG
- Transportation
- Insurance
- Telecommunications
- Other

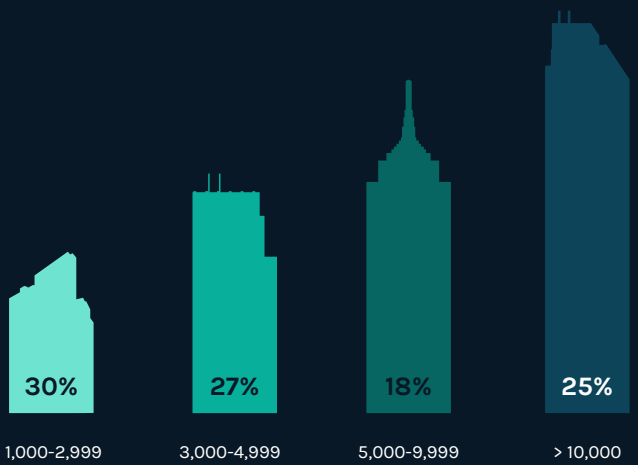

Number of employees in organization

Survey Respondent Demographics

1,035 U.S.-based respondents were surveyed to explore Generative AI and AI trends and challenges.

Primary Role

31% Platform Team/Engineering/Operations

1% Platform Architecture

18% IT/Cloud Architecture

23% IT/Cloud Engineering

19% Developer/Engineer

6% DevOps

2% SRE/Operations

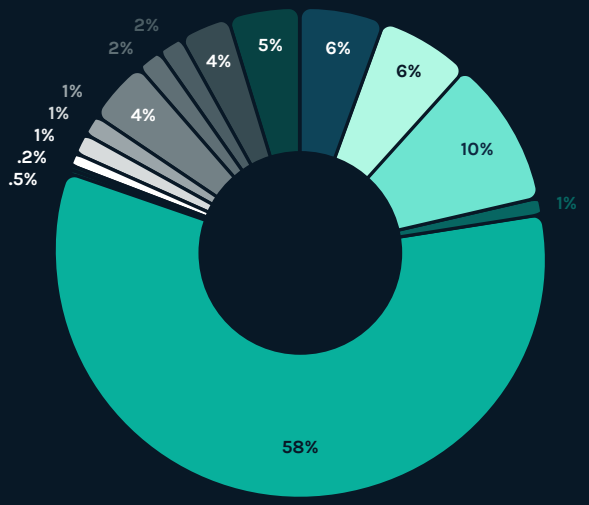

Industry

- Healthcare
- Financial Services
- Manufacturing
- Education
- Computer Software/Hardware
- Media & Entertainment
- Marketing & Advertising
- Hospitality
- Public Administration/Government
- Energy
- Retail & CPG
- Transportation
- Insurance
- Telecommunications
- Other

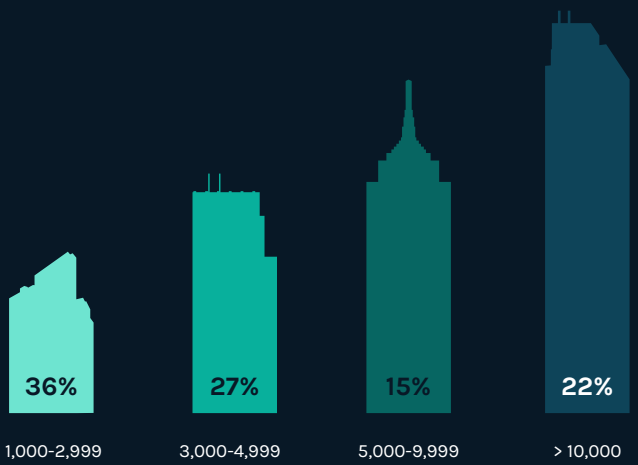

Number of employees in organization

Learn more at **rafay.co**