STATE OF CLOUD NATIVE DEVELOPMENT Q1 2025

/DATA

CLOUD NATIVE COMPUTING FOUNDATION

MARCH 2025

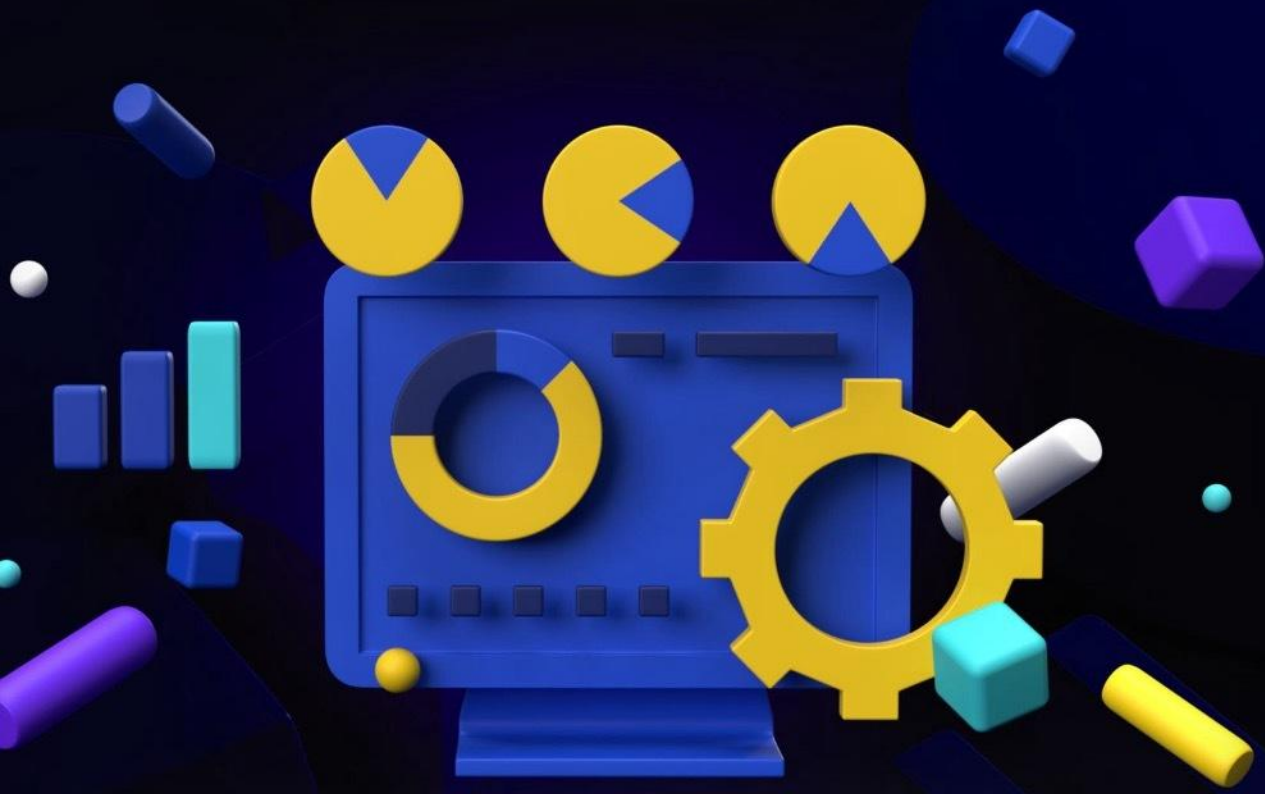

We transform noise into actionable, high-value insights for decision-makers

SlashData is a Market Research firm, serving the technology and more industries for 20+ years. We support the strategic decision-making of industry leaders with exceptional quality data. With our help, CEOs, CTOs, strategic management teams, product managers and engagement and support professionals have the confidence and information they need to invest in the right audiences and markets, maximising ROI and driving the industry forward to the future.

From product development to audience insights and competitive intelligence, we can get you the answers you need.

www.slashdata.co

Can I share data from this report?

1. License Grant
This report is licensed under the Creative Commons Attribution-NoDerivatives Licence 4.0 (International) . Put simply, subject to the terms and conditions of this license, you are free to:

Share — You can reproduce the report or incorporate parts of the report into one or more documents or publications, for commercial and non-commercial purposes.

Under the following conditions:

Attribution — You must give appropriate credit to SlashData™, and to the Continuous Delivery Foundation as sponsors of this report, and indicate if changes were made. In that case, you may do so in any reasonable manner, but not in any way that suggests that SlashData™ endorses you or your use.

NoDerivatives — you cannot remix or transform the content of the report. You may not distribute modified content.

2. Limitation of Liability
SlashData™, believes the statements contained in this publication to be based upon information that we consider reliable, but we do not represent that it is accurate or complete and it should not be relied upon as such. Opinions expressed are current opinions as of the date appearing in this publication only and the information, including the opinions contained herein, are subject to change without notice. Use of this publication by any third party for whatever purpose should not and does not absolve such third party from using due diligence in verifying the publication's contents. SlashData™ disclaims all implied warranties, including, without limitation, warranties of merchantability or fitness for a particular purpose.

SlashData™, its affiliates, and representatives shall have no liability for any direct, incidental, special, or consequential damages or lost profits, if any, suffered by any third party as a result of decisions made, or not made, or actions taken, or not taken, based on this publication.

The analyst of the developer economy | formerly known as VisionMobile

ABOUT THE AUTHORS

Liam Bollmann – Dodd

Senior Market Research Consultant

Liam is a former experimental antimatter physicist, and he obtained a PhD in Physics while working at CERN. He is interested in the changing landscape of cloud development, cybersecurity, and the relationship between technological developments and their impact on society.

liam.dodd@slashdata.co

Nikita Solodkov

Market Research and Statistics Consultant

Nikita Solodkov is a multidisciplinary researcher with a particular interest in using data-driven insights to solve real-world problems. He holds a PhD in Physics and has over five years of experience in data analytics and research design.

nikita.solodkov@slashdata.co

TABLE OF CONTENTS

KEY INSIGHTS

- In Q1 2025, an estimated 49% of backend service developers are cloud native. This accounts for approximately 9.2M developers. →
- 93% of all developers are deploying to cloud for some part of their development process. This rises to 96% for those in backend service development. →
- Cloud functions or serverless architectures are currently used by 20% of backend service developers, down from 26% in Q3 2020. →
- The proportion of backend service developers using containers has remained relatively consistent for the last four and a half years, at 61%. →
- The majority of professional developers in backend services, machine learning (ML), artificial intelligence (AI), or data science are not using Kubernetes for their ML/AI workloads. →
- 18% of these developers are planning to use Kubernetes for their ML/AI workloads. →
- More than half of professional ML developers and data scientists are using Kubernetes for AI/ML workloads. →
- Developers working on model infrastructure and pipeline engineering or developing ML frameworks/libraries are the most likely to be using Kubernetes for AI/ML workloads. →

INTRODUCTION

01

1. Introduction

Cloud deployment has become an integral part of software development, with the vast majority of developers now relying on cloud infrastructure in some capacity. Despite this ubiquity, not all developers are fully leveraging the capabilities of cloud computing through cloud native technologies and approaches.

This report, produced in partnership with the Cloud Native Computing Foundation (CNCF), presents the latest insights into the state of cloud native development in Q1 2025. This report provides an analysis of key trends shaping the cloud native ecosystem, drawing on data from the 29th edition of SlashData's Developer Nation survey, fielded between December 2024 and February 2025, and reaching more than 10,000 developers worldwide.

For this report, cloud native development is determined by the use of a combination of technologies, including containers, container orchestration and management tools, service meshes, Kubernetes, and cloud functions or serverless computing. By tracking adoption patterns and technology usage, this report offers a comprehensive update on the growth of cloud native development, the evolving developer landscape, and shifts in cloud deployment strategies.

1. Introduction

A key focus of this year's report is the role of Kubernetes in supporting artificial intelligence (AI) / machine learning (ML) workloads. We examine the prevalence of Kubernetes in AI/ML pipelines, the types of workloads most commonly run on Kubernetes, and the specific AI/ML activities that are leading adoption.

By analysing these trends, this report provides a clearer picture of how cloud native technologies are shaping modern software development and where the industry is headed next.

TRENDS IN CLOUD NATIVE DEVELOPMENT

02

2. Trends in Cloud Native Development

In Q1 2025, the proportion of developers in backend services who we identify as cloud-native is 49%. This is an increase from Q1 2021, where we estimated 46% of backend developers were cloud-native. While the proportion of cloud native developers has only increased slightly, the absolute number of developers has increased substantially due to the growth in developer numbers. As of Q1 2025, we estimate that there are approximately 9.2M cloud-native backend developers[1]. This corresponds to an increase of 2.3M from Q1 2021 (6.9M) and 0.5M from Q3 2023 (8.7M).

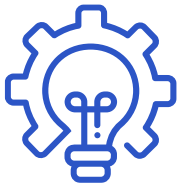

9.2M

CLOUD NATIVE BACKEND DEVELOPERS IN Q1 2025

[1] *The 95% confidence interval for this estimate is given by (8.8M, 9.6M).*

2. Trends in Cloud Native Development

In addition to backend services, cloud-native practices among those in ML and AI are also of growing interest. As AI/ML workloads become more complex and resource-intensive, cloud native technologies offer scalability, automation, and efficiency benefits that are critical for modern AI development. As of Q3 2024 we estimate that 24% of developers in machine learning, AI, or data science (MLDS) are cloud native.

While our historical data on this group is reduced, this is a 5 percentage point increase from Q1 2024, suggesting that cloud native practices may be increasing among MLDS developers. This is further supported by our data on Kubernetes usage, shown in chapter 3. We estimate that there were 4.2M MLDS developers in Q3 2024.

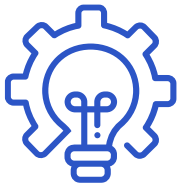

4.2M

CLOUD NATIVE MLDS DEVELOPERS IN Q3 2024

The proportion of cloud native developers among backend services and MLDS developers

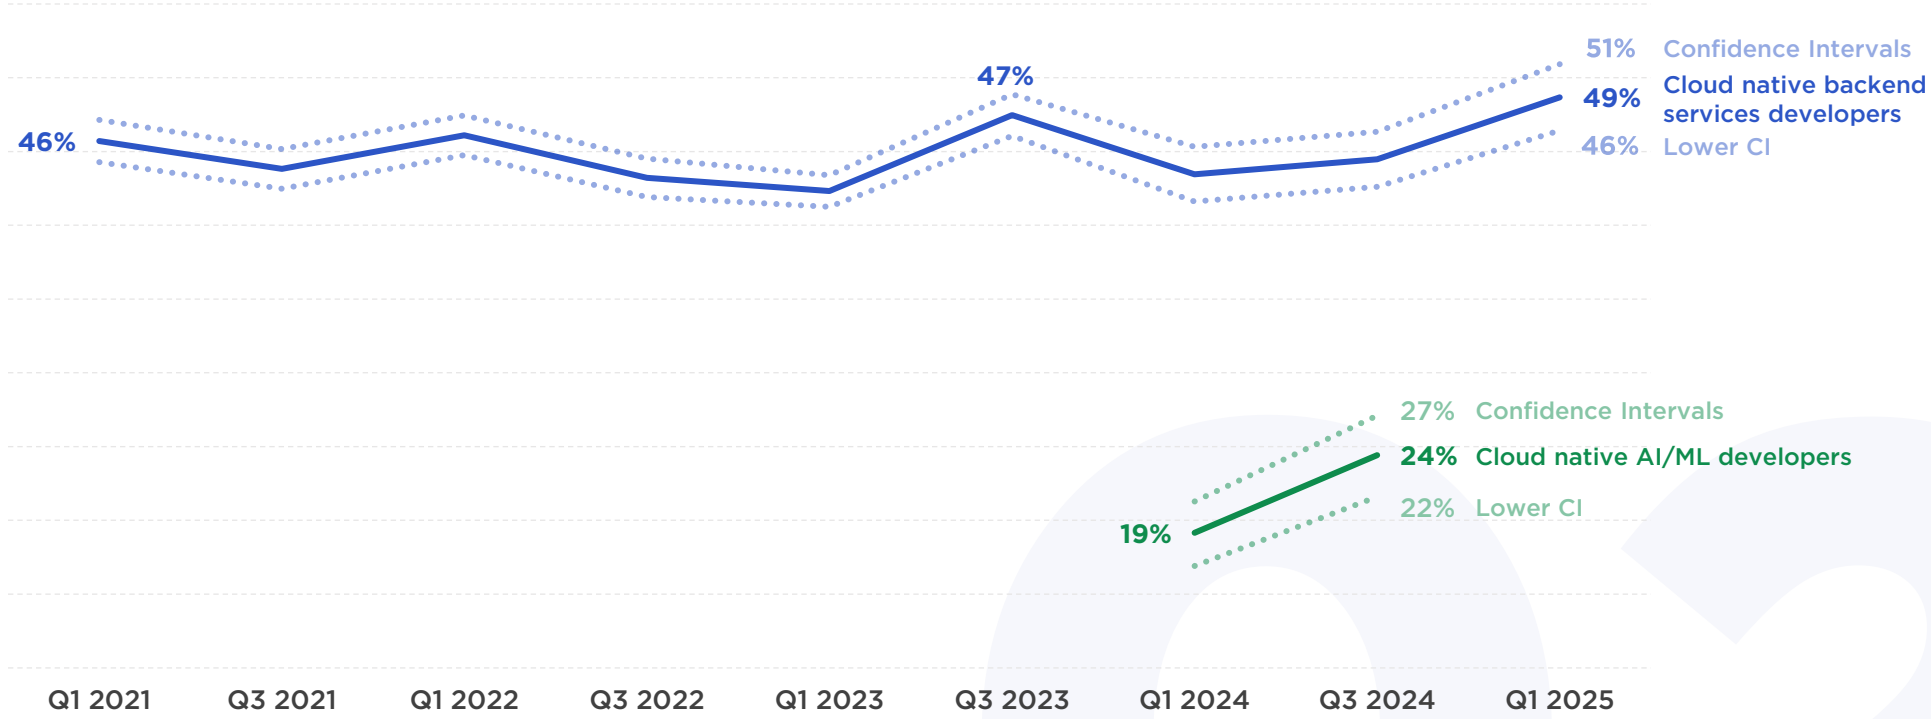

Question wording: Which of the following technologies have you used as part of your backend services development in the last 12 months? | Which of the following technologies have you used for data science / machine learning in the last 12 months?
% of developers by development area (21Q1 n=4,668 | 21Q3 n=5,260 | 22Q1 n=5,261 | 22Q3 n=5,619 | 23Q1 n=8,182 | 23Q3 n=4,847 | 24Q1 n=2,743 | 24Q3 n=2,754 | 25Q1 n=1,903)

2. Trends in Cloud Native Development

In addition to trends in cloud native developers, we can also examine how cloud deployment has changed over time. When looking at all developers, the usage of hybrid cloud has increased to 29% from 22% in Q1 2021. Multicloud has seen a similar increase over the same period, growing from 17% of all developers to 22%. The story is different for public and private clouds. Private cloud has only increased by 2 percentage points over the last four years, reaching 27% in Q1 2025, while public cloud has decreased by 2 percentage points, to 36%.

When paired with the increased usage of on-premises servers and mainframes, it may appear that cloud deployment is stationary or losing ground. However, 93% of developers are deploying to the cloud in some capacity, even if in combination with non-cloud options. This is an increase from 85% in Q1 2021 and 87% in Q1 2023.

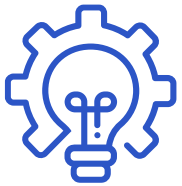

93%

OF DEVELOPERS DEPLOY TO THE CLOUD

2. Trends in Cloud Native Development

Among backend service developers, we see similar trends but with more variation over the same period. Hybrid and multicloud (23% and 16%, respectively) have both increased as a proportion of deployment options since Q1 2021 (20% and 13%, respectively), but hybrid cloud is a smaller proportion than its peak of 27% in Q1 2023.

Public and private cloud are currently at their lowest shares of popularity since Q1 2021 (33% each). However, 96% of backend service developers are currently using cloud in some capacity for their projects and deployments, up from 86% in Q1 2023. Further, this does not measure how much compute they are using the cloud for, simply the diversity in their compute types.

The reduction of private and public cloud usage, as a proportion, is a combination of two factors. Firstly, backend service developers are consolidating around fewer cloud approaches, despite their increasing usage of cloud overall. Secondly, the introduction of a new option, distributed cloud[1] (12%, Q1 2025), has impacted the selection of other options as developers find a more accurate identification for their deployment approach.

[1] *This refers to geographically distributed compute that is managed centrally.*

Trends in cloud deployment

Cloud deployment trends among all developers

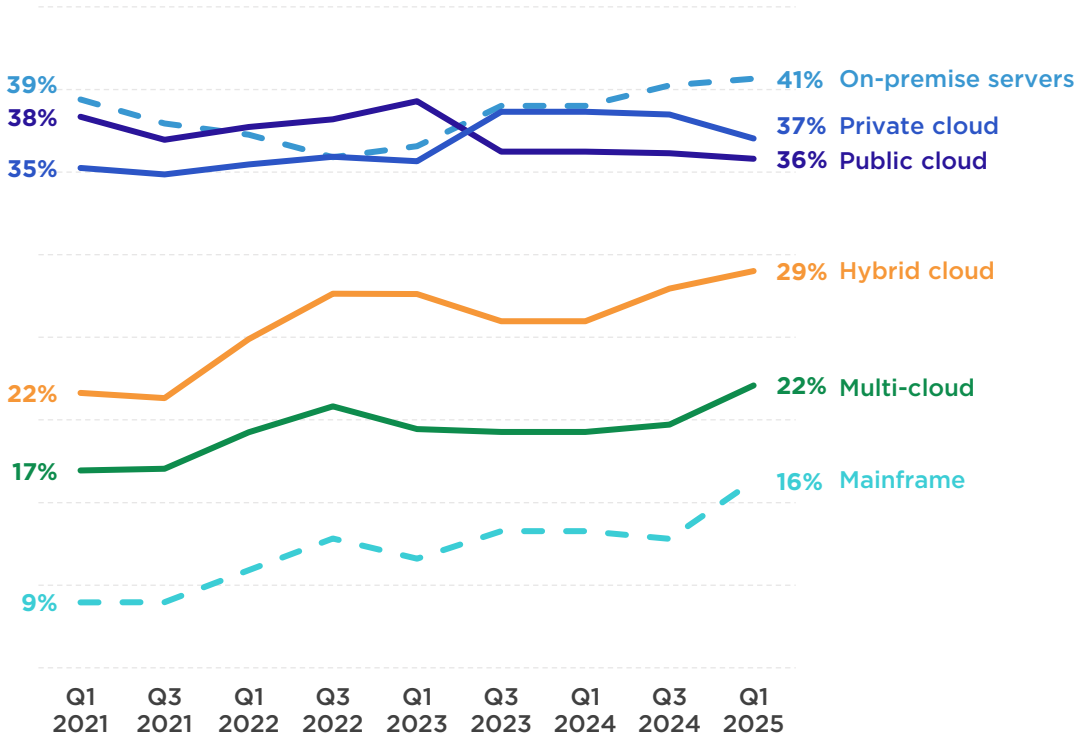

Trends in deployment of backend service developers

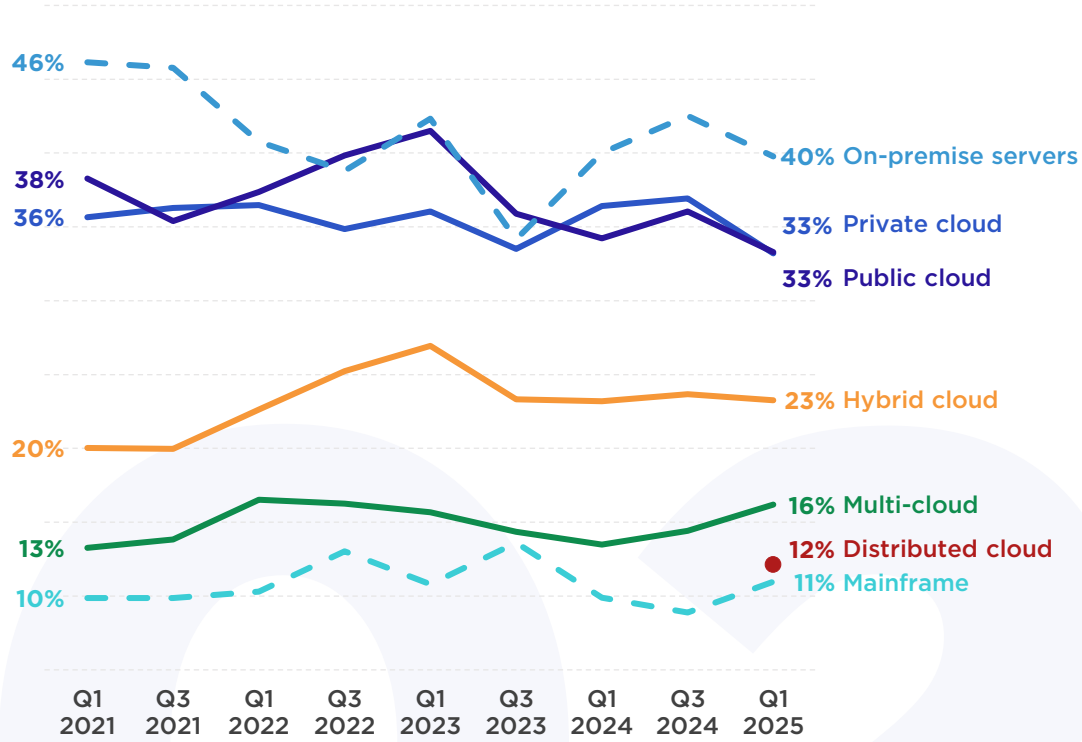

Question wording: Where does the code of your backend services run?
% of developers | % of backend service developers (21Q1 n=11,177 | 21Q3 n=12,132 | 22Q1 n=13,344 | 22Q3 n=17,820 | 23Q1 n=20,879 | 23Q3 n=10,521 | 24Q1 n=10,521 | 24Q1 n=9,886 | 25Q1 n=10,731)

2. Trends in Cloud Native Development

Cloud native development is not solely about the technologies used. However, tracking the usage of technologies associated with cloud native development provides key context to how the space is developing.

The usage rate of containers by backend service developers has remained relatively consistent over the last four years (61% in Q1 2025). Meanwhile, Kubernetes has shown a small positive trend in its adoption (31% in Q1 2025). The absolute adoption of these technologies has grown alongside the growth in the backend service developer population. The consistent usage rates may indicate that these are mature technologies, which won't see any major spikes or drops in usage rates without other major changes in the backend space.

The usage rate of containers (61% in Q1 2025) has remained stable for four years.

Container orchestration tools and cloud functions / serverless architecture have seen their usage rates decrease substantially from Q3 2023. Container orchestration has decreased from 29% in Q3 2023 to 22% in Q1 2025, and cloud functions and serverless architecture have decreased from 26% to 20% over the same period.

The reduction in developers stating their usage of container orchestration tools may be a consequence of cloud service providers abstracting away or integrating orchestration within their workflows. This reduces both the explicit need for these tools and blurs the lines for whether developers are using such tools.

The reduction in the usage rate of cloud functions and serverless architecture may have multiple drivers. However, if backend service developers are supporting more applications and services involved in AI or ML, then these technologies are less likely to be cost-effective. Serverless architecture works well for intermittent services that scale up and down based on needs, which are less useful for many high-usage tasks, such as ML inference.

Trends in adoption of technologies associated with cloud-native development

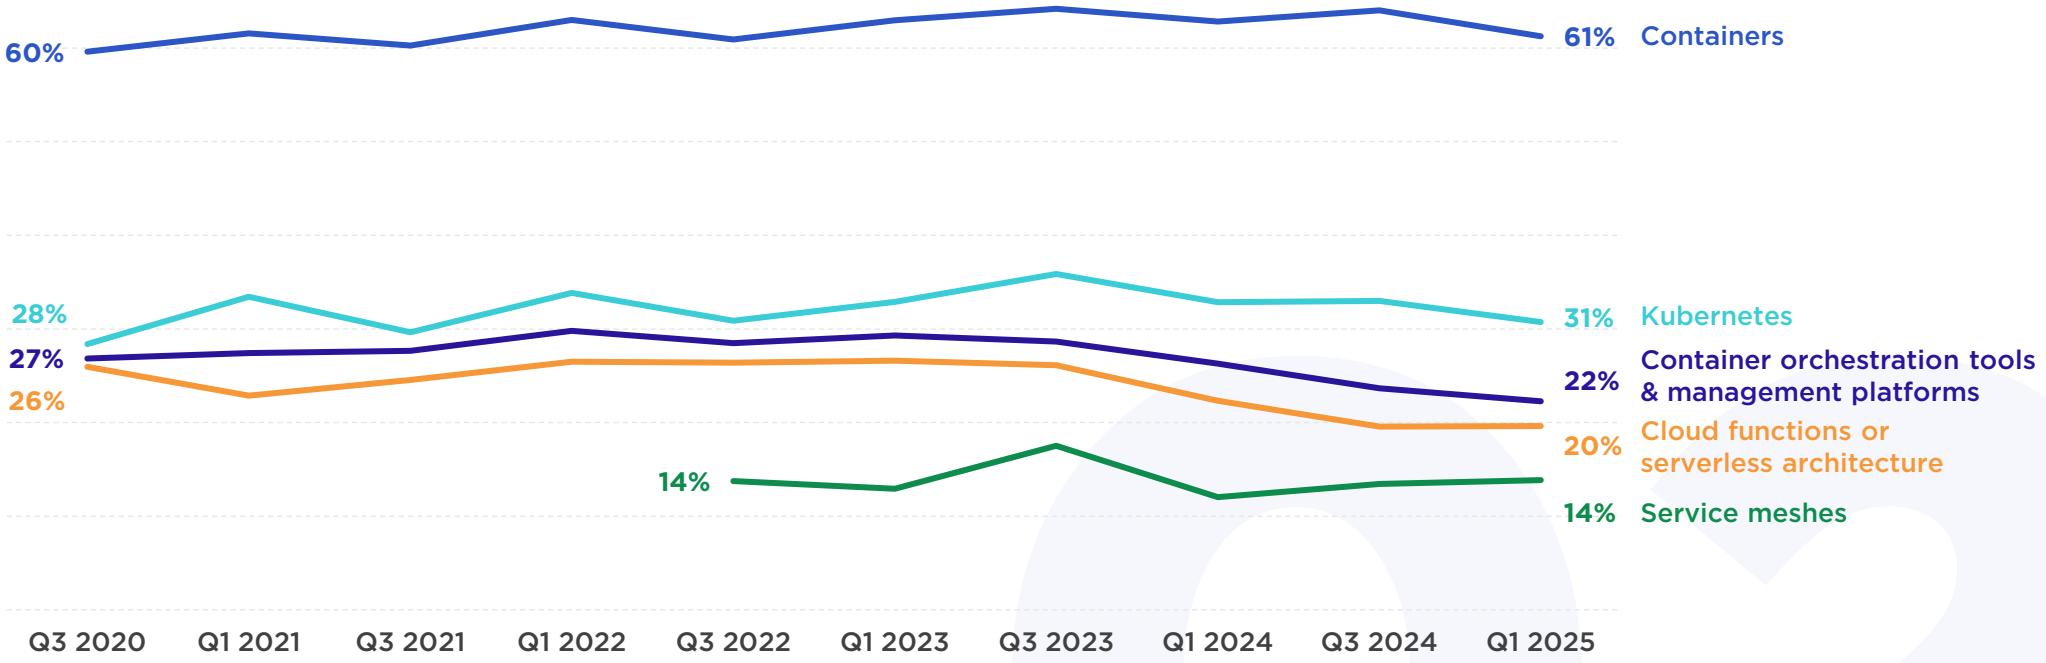

Question wording: Which of the following technologies have you used as part of your backend services development in the last 12 months?
% of backend service developers (21Q1 n=2,703 | 21Q3 n=2,798 | 22Q1 n=3,166 | 22Q3 n=4,806 | 23Q1 n=7,877 | 23Q3 n=2,781 | 24Q1 n=3,430 | 24Q1 n=3,226 | 25Q1 n=3,621)

KUBERNETES IN ML/AI WORKLOADS

03

3. Kubernetes in ML/AI Workloads

We asked developers involved professionally in backend services, ML/AI, or data science whether they use Kubernetes for ML/AI workloads. If they did, we asked them which specific workloads they were running on Kubernetes. The majority (64%) of these developers were ***not*** using Kubernetes for ML/AI workloads. However, 18% of these developers indicated that they were planning to do so in the future.

36% of professional MLDS or backend developers are running ML/AI workloads on Kubernetes, with a further 18% more planning to

3. Kubernetes in ML/AI Workloads

The most common ML/AI workloads on Kubernetes were data pre-processing (11%), model experimentation (9%), and real-time model inference (9%). The least popular workloads were feature engineering (7%) and hyperparameter tuning (6%). However, the lower proportion of hyperparameter tuning is likely due to its niche role within ML/AI workflows rather than any fundamental limitation of running it on Kubernetes.

While Kubernetes can technically support a wide range of ML/AI activities, some tasks may be better suited to alternative environments. For instance, large-scale batch data processing often benefits from specialised big data frameworks. The relative lack of popularity of ML/AI workloads on Kubernetes likely reflects the overall prevalence of these tasks rather than any inherent limitation of running them on Kubernetes.

ML/AI workloads ran on Kubernetes

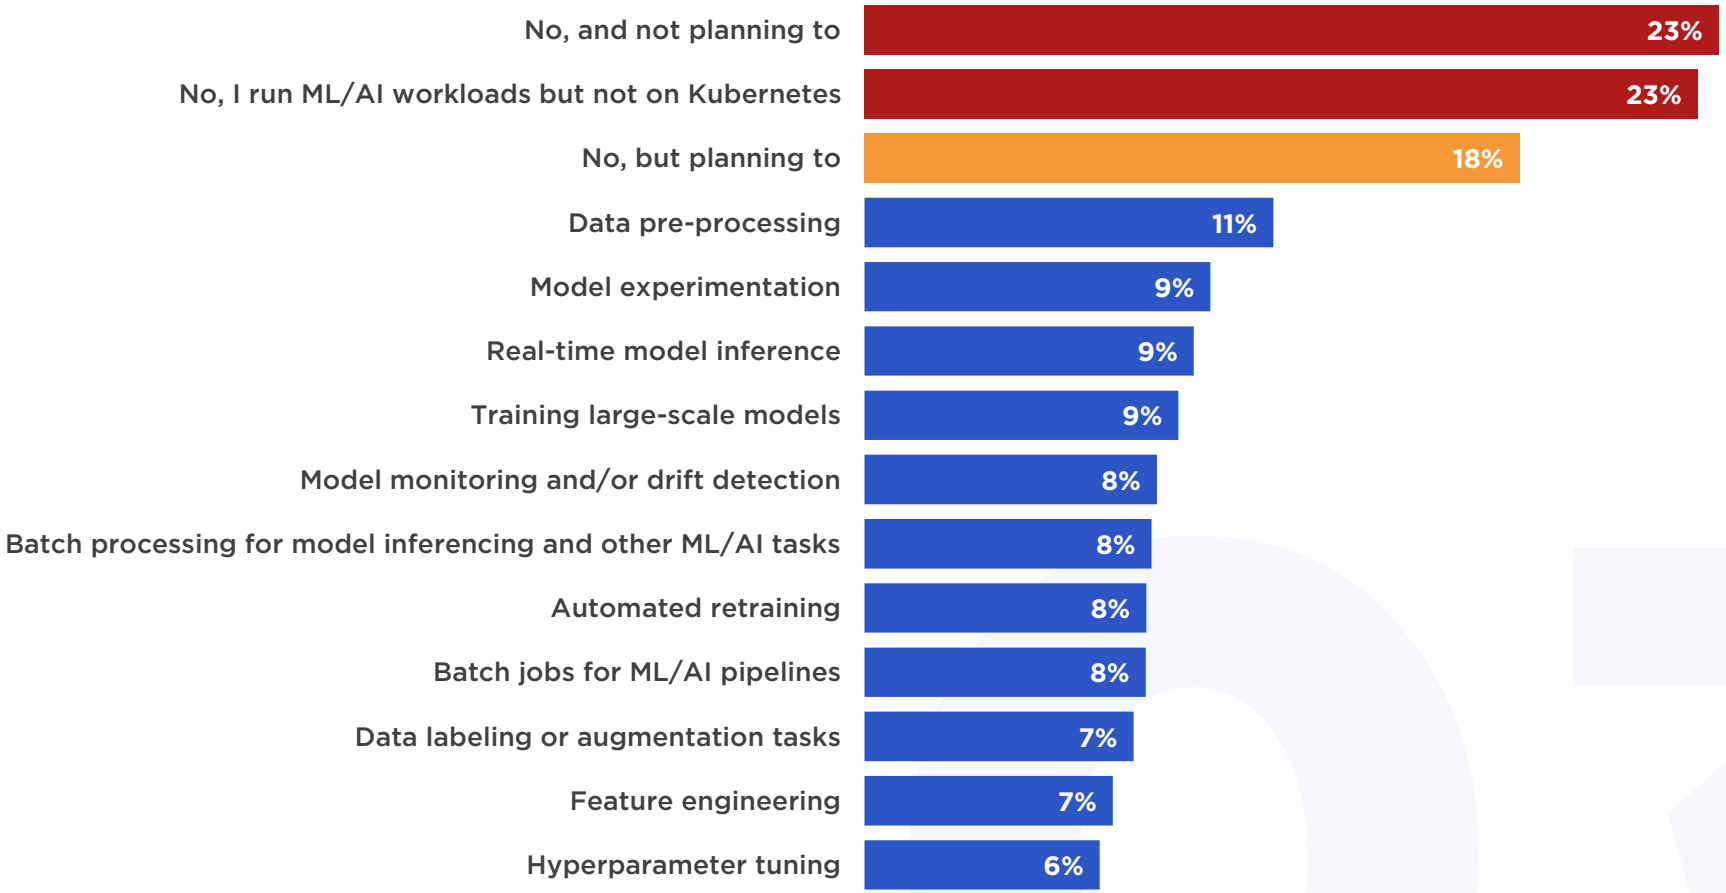

Question wording: Are you or your team running ML/AI workloads on Kubernetes? If so, which types of tasks and workloads do you run on Kubernetes?
% of developers involved professionally in backend development, ML/AI, or data science (Q1 2025 n=4,456)

3. Kubernetes in ML/AI Workloads

While many types of developers may be involved in running ML/AI workloads on Kubernetes, it is significantly more commonplace for some than others. Developers were asked to self-identify with the kind of development projects they work on, which we can use to categorise them into groups.

Data scientists and ML/AI developers are substantially more likely to be running ML/AI workloads on Kubernetes (50% and 52%, respectively). As these developers are primarily focused on ML/AI, this result is unsurprising. However, it does highlight that the skills expected, or required, of MLDS developers include various aspects of backend infrastructure too.

Modern professional MLDS developers are expected to have some proficiency with various aspects of backend infrastructure

3. Kubernetes in ML/AI Workloads

38% of DevOps developers who are involved professionally in backend services, ML/AI, or data science are running ML/AI workloads on Kubernetes. Despite the broad range of responsibilities that fall under DevOps, this highlights the critical role of reliable DevOps practices in ML/AI development. DevOps developers working with ML/AI workloads on Kubernetes are more likely to be involved in model experimentation (32%) and real-time model inference (28%) compared to other developer groups. This could be due to their involvement in managing infrastructure and optimising CI/CD pipelines.

Among developers using Kubernetes for ML/AI workloads, we observe a consistent hierarchy of task popularity across different groups. This supports the hypothesis that task popularity on Kubernetes is primarily a reflection of overall ML/AI task prevalence rather than technical suitability.

While ML/AI developers and data scientists have similar Kubernetes adoption rates, their usage patterns show some distinctions. Data scientists are more likely to use Kubernetes for data pre-processing (35% vs 32%), reflecting their focus on preparing and transforming data before model training. Meanwhile, ML/AI developers are more likely to use Kubernetes for batch processing of model inference and other tasks (24% vs 20%) and automated retraining (25% vs 21%). This aligns with their role in scaling and automating model deployments, ensuring that production models remain up-to-date and performant.

Usage of Kubernetes for ML/AI workloads by development area involved in

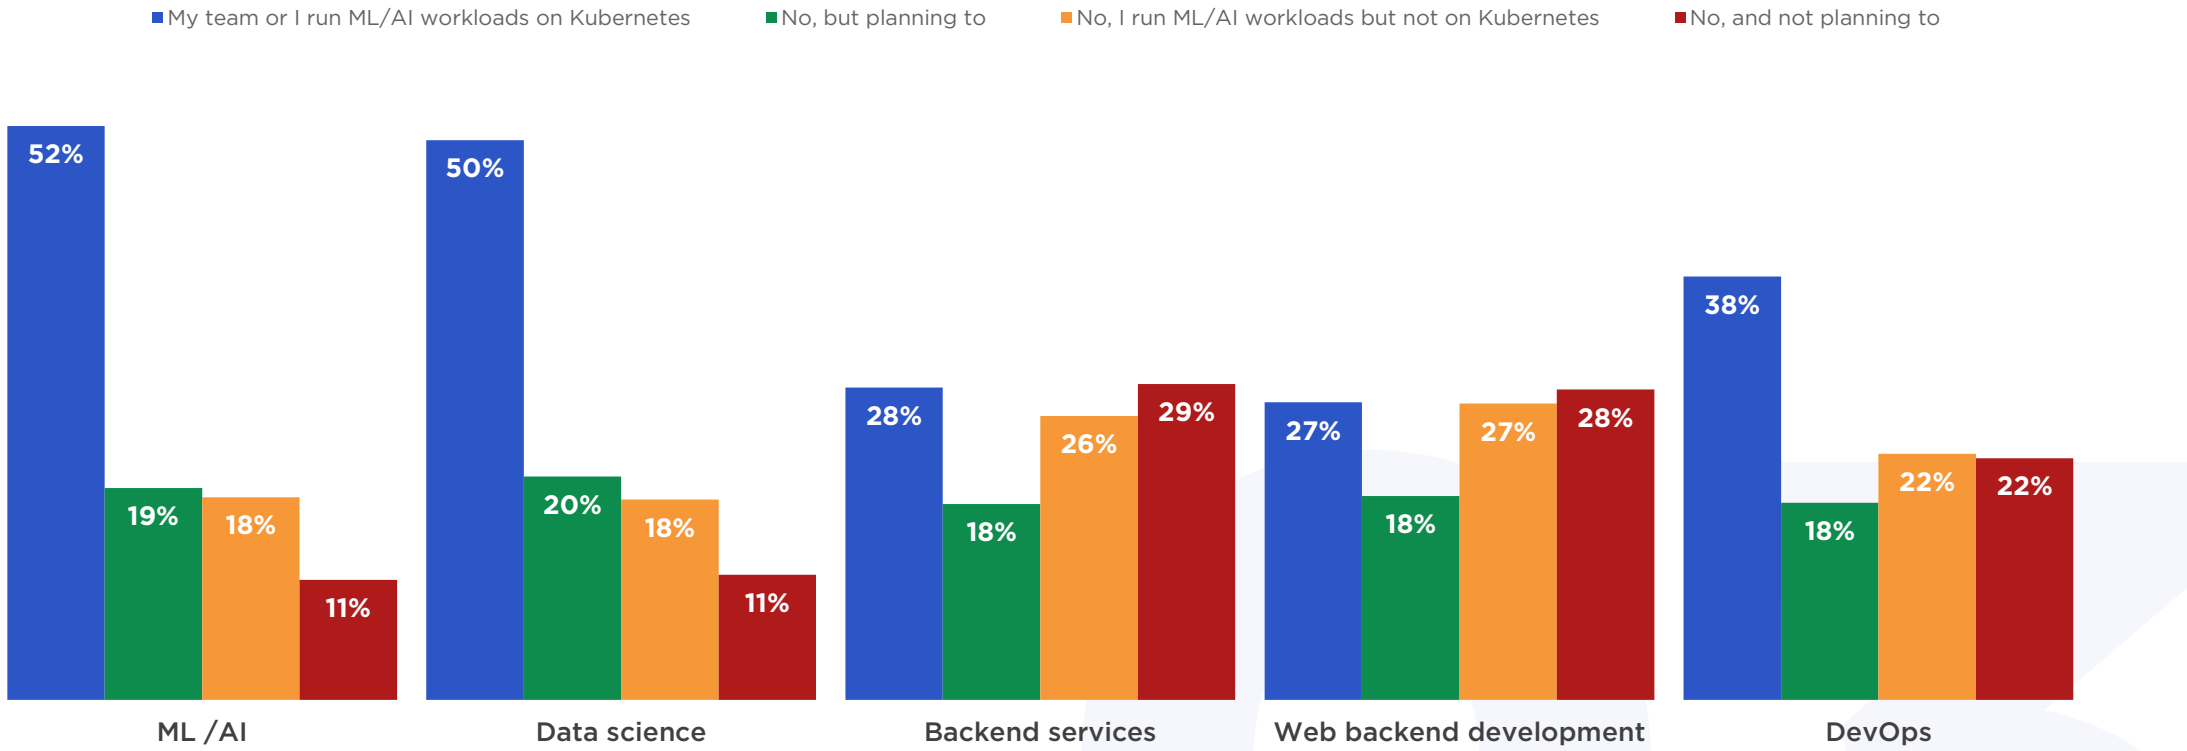

Question wording: Which of the following types of apps / development projects are you involved in? | Are you or your team running ML/AI workloads on Kubernetes? If so, which types of tasks and workloads do you run on Kubernetes?
% of developers who are involved professionally in backend services, ML/AI, or data science by development area (ML/AI n=1,825 | Data science n=1,393 | Backend services n=3,205 | Web backend development n=2,273 | DevOps n=1,149)

ML/AI workloads ran on Kubernetes by development area involved in

Areas of development involved in

Tasks on Kubernetes

	ML/AI	Data science	Backend services	Web backend development	DevOps
Data pre-processing	32%	35%	30%	32%	33%
Model experimentation	27%	27%	28%	30%	32%
Real-time model inference	25%	25%	24%	27%	28%
Training large-scale models	25%	24%	23%	24%	23%
Model monitoring and/or drift detection	23%	23%	22%	23%	19%
Batch processing for model inferencing and other ML/AI tasks	24%	20%	20%	21%	21%
Automated retraining	25%	21%	20%	24%	24%
Batch jobs for ML/AI pipelines	22%	24%	19%	23%	22%
Data labeling or augmentation tasks	23%	23%	17%	20%	21%
Feature engineering	20%	22%	17%	20%	22%
Hyperparameter tuning	18%	19%	15%	14%	18%

- >5pp below the average of all other cells within each row
- 2.5 – 5pp below the average of all other cells within each row
- ±2.5pp around the average of all other cells within each row
- 2.5 – 5pp above the average of all other cells within each row
- >5pp above the average of all other cells within each row

Question wording: Which of the following types of apps / development projects are you involved in? | Are you or your team running ML/AI workloads on Kubernetes? If so, which types of tasks and workloads do you run on Kubernetes?
% of developers who are involved professionally in backend services, ML/AI, or data science and run ML/AI workloads on Kubernetes by development area (ML/AI n=945 | Data science n=703 | Backend services n=904 | Web backend development n=610 | DevOps n=439)

3. Kubernetes in ML/AI Workloads

In the final section, we focus on MLDS developers and break down this information by the types of MLDS activities that they identified involvement with. Here, we examine which workloads are most commonly run on Kubernetes, as well as their likelihood to use Kubernetes for ML/AI workloads.

Amongst MLDS developers, the least likely to use Kubernetes for ML/AI workloads are those who use third-party models through APIs (48%), engage in data analysis, exploration, or reporting (46%), or focus on model training (47%). These tasks often do not require the infrastructure orchestration that Kubernetes provides, as they can be performed efficiently on managed services or local environments. Conversely, those involved in model infrastructure and pipeline engineering (71%), developing ML frameworks or libraries (70%), and teaching ML, AI, or data science (69%) are the most likely to leverage Kubernetes, reflecting the need for scalable, automated, and repeatable workflows in these roles.

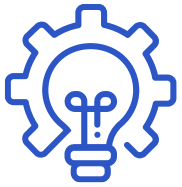

71%

OF MLDS DEVELOPERS WORKING ON MODEL INFRASTRUCTURE OR PIPELINE ENGINEERING ARE USING KUBERNETES FOR ML/AI WORKLOADS

3. Kubernetes in ML/AI Workloads

Alongside their high Kubernetes adoption, developers working on ML frameworks and model infrastructure demonstrate above-average Kubernetes usage across a wide range of workloads. Both groups show the highest Kubernetes usage for real-time model inference (25% and 24%, respectively), while those working on model infrastructure also lead in feature engineering participation (19%), suggesting a focus on optimising data transformations and model inputs at scale.

ML/AI workloads ran on Kubernetes by development area involved in

ML/AI activities

ML/AI workloads on Kubernetes	Conducting ML, AI, or data science research	Studying ML, AI, or data science	Teaching ML, AI, or data science	Training models	Deploying models	Data analysis, exploration or reporting	Data engineering	Using third-party models through APIs	Developing ML frameworks or libraries	Model infrastructure/pipeline engineering	Translating business problems into data problems
Not using Kubernetes for ML/AI workloads	39%	45%	31%	47%	40%	46%	44%	48%	30%	29%	44%
Data pre-processing	21%	20%	23%	18%	22%	21%	21%	19%	25%	21%	18%
Real-time model inference	18%	15%	20%	18%	19%	18%	17%	17%	25%	24%	17%
Training large-scale models	18%	16%	17%	16%	16%	14%	13%	15%	21%	21%	16%
Data labeling or augmentation tasks	18%	16%	20%	13%	15%	16%	15%	15%	18%	18%	16%
Model experimentation	16%	15%	17%	17%	16%	15%	16%	13%	16%	20%	16%
Automated retraining	16%	15%	18%	15%	17%	14%	13%	15%	18%	20%	15%
Batch processing for model inferencing and other ML/AI tasks	18%	14%	17%	12%	12%	14%	12%	11%	18%	16%	11%
Batch jobs for ML/AI pipelines	16%	13%	21%	13%	16%	14%	14%	15%	18%	20%	13%
Model monitoring and/or drift detection	16%	13%	20%	14%	15%	12%	15%	13%	19%	18%	15%
Feature engineering	14%	12%	14%	11%	13%	12%	13%	9%	15%	19%	14%
Hyperparameter tuning	14%	11%	13%	10%	12%	10%	12%	10%	17%	16%	11%

- >5pp below the average of all other cells within each row
- 2.5 – 5pp below the average of all other cells within each row
- ±2.5pp around the average of all other cells within each row
- 2.5 – 5pp above the average of all other cells within each row
- >5pp above the average of all other cells within each row

Question wording: Which of the following types of apps / development projects are you involved in? | Are you or your team running ML/AI workloads on Kubernetes? If so, which types of tasks and workloads do you run on Kubernetes?
% of developers who are involved professionally in backend services, ML/AI, or data science and run ML/AI workloads on Kubernetes by development area (ML/AI n=945 | Data science n=703 | Backend services n=904 | Web backend development n=610 | DevOps n=439)

METHODOLOGY

The Developer Nation Survey

GLOBAL REACH

Geo distribution of respondents of the 29th global Developer Nation Survey (December 2024 to February 2025)

29th
GLOBAL WAVE

10,500+
DEVELOPERS

126
COUNTRIES REACHED

Developer Nation 29th edition reached 10,500+ respondents around the world. As such, the Developer Nation series continues to be the most comprehensive independent research on mobile, desktop, Industrial IoT, consumer electronics, 3rd party app ecosystems, cloud, web, game, AR/VR and machine learning developers and data scientists combined ever conducted. The report is based on the large-scale online developer survey designed, produced and carried out by SlashData over a period of 11 weeks between December 2024 and February 2025.

Respondents to the online survey came from 125+ countries, including the US, China, India, Israel and the UK. The geographic reach of this survey is reflective of the global scale of the developer economy. The online survey was translated into nine languages in addition to English (Simplified Chinese, Traditional Chinese, French, Spanish, Portuguese, Vietnamese, Russian, Japanese, Korean) and promoted by more than 70 leading community and media partners within the software development industry.

Our respondents came from a broad age spectrum, from young coders and creators who are under 18 to the seasoned ones over 55. Excluding those who would rather not answer about their age, the age profile of our respondents is shown below.

Age distribution of survey respondents

% of respondents (Q1 2025 n=10,532)

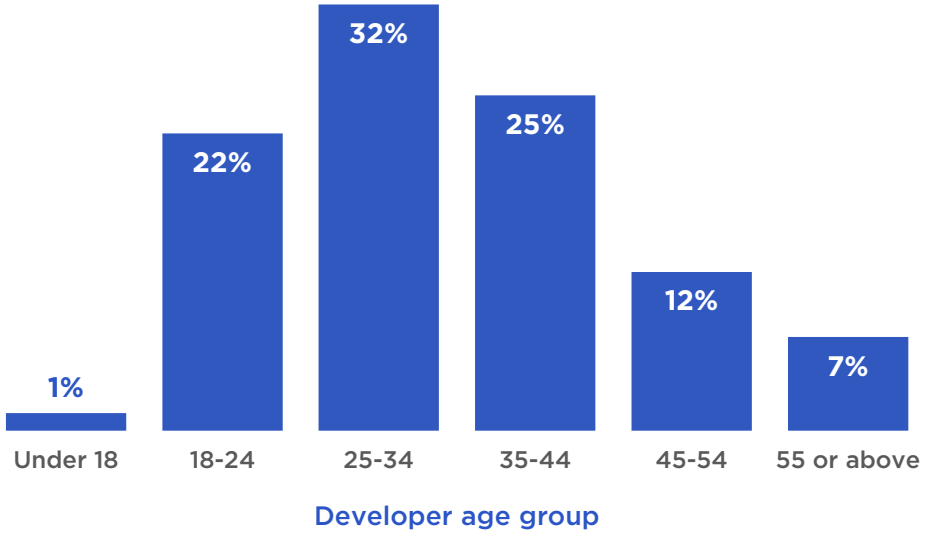

How respondents are involved in each software area

% of respondents in each area (Q1 2025 n=10,532)

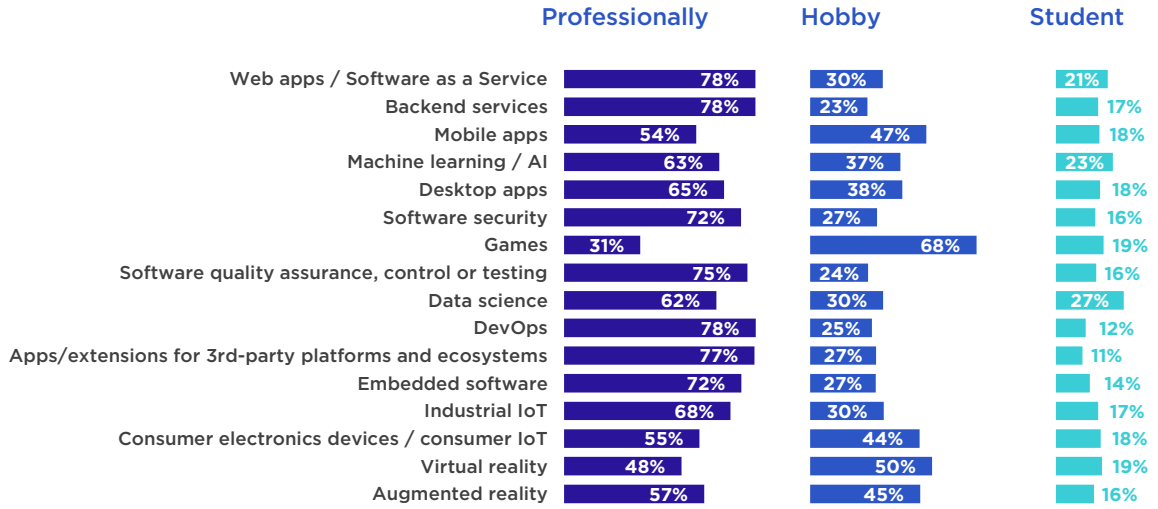

Respondents were asked which types of projects they are involved in out of the 13 under study, namely web apps / SaaS, mobile apps, desktop apps, backend services, augmented reality, virtual reality, games, data science, machine learning / artificial intelligence, industrial IoT, consumer electronics devices, embedded software, and apps/extensions for third-party app ecosystems. They also told us if they are into their areas of involvement as professionals, hobbyists, or students - or as any combination of these - and how many years of experience they have in each.

METHODOLOGY

To eliminate the effect of regional sampling biases, we weighted the regional distribution across nine regions by a factor that was determined by the regional distribution and growth trends identified in our Developer Nation research. To minimise other important sampling biases across our outreach channels, we weighted the responses to derive a representative distribution for technologies used, and developer segments. Using ensemble modeling methods, we derived a weighted distribution based on data from independent, representative channels, excluding the channels of our research partners to eliminate sampling bias due to respondents recruited via these channels. Each of the separate branches: Industrial IoT, consumer electronics, 3rd party app ecosystems, cloud, embedded, augmented and virtual reality were weighted independently and then combined.

For more information on our methodology please visit https://www.slashdata.co/methodology.

TRUSTED BY

the leading tech platforms

OUR CLIENTS


/DATA
CLOUD NATIVE
COMPUTING FOUNDATION