CNCF TECHNOLOGY LANDSCAPE RADAR

MULTICLUSTER APPLICATION MANAGEMENT AND BATCH/AI/ML COMPUTE TECHNOLOGIES

/DATA

CLOUD NATIVE COMPUTING FOUNDATION

NOVEMBER 2024

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 Dodd

Senior Market Research Analyst

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

Brayton Noll

Senior Market Research Analyst

Brayton Noll is a behavioral scientist with a background in climate change and environmental research. He holds a PhD from TU Delft in computational social-science with his thesis focusing on human behavioral dynamics and climate adaptation. He has five years of experience working with data analytics.

brayton.noll@slashdata.co

TABLE OF CONTENTS

KEY INSIGHTS

- For multicluster application management technologies, ArgoCD and Cilium were considered technologies ready to ‘adopt’. →
- Cilium was considered the most mature and useful technology of this group. →
- ArgoCD was also considered highly useful, but ranked lower on the maturity rankings, but was the multicluster application management technology that received the highest proportion of developers who would recommend the technology. →
- For Batch/AI/ML computing technologies, Apache Airflow, CubeFS, Kubeflow, and Fluid were placed in the ‘adopt’ position of the technology radar. →
- Apache Airflow, CubeFS, and Kubeflow were the top three ranked technologies for usefulness and maturity. →
- Fluid ranked highly on maturity, and was one of the most likely technologies developers stated they would recommend to others. →

INTRODUCTION

01

1. Introduction

In Q3 2024, more than 300 professional developers using technologies associated with cloud-native development were asked about their experience and opinions with regard to batch computing and multicluster application management technologies.[1] The developers originate from around the world, and have a large range of specialities and areas of focus. A more granular breakdown of the respondents is included in the Methodology section.

For the products or tools they were familiar with, they rated them on their usefulness and maturity and indicated how likely they were to recommend that technology to other developers. Within the context of this report, usefulness was defined as how well it meets project requirements, and maturity related to a technologies' stability and reliability. The recommendation scale was converted into a net promoter score (NPS) for use during the analysis.

Based on the usage, usefulness and maturity ratings, and how likely they are to recommend a given technology, we categorised the technologies into three groups: adopt, trial and assess. 'Adopt' technologies are considered reliable choices for most use cases, while 'trial' technologies are worth exploring to see if they meet your specific needs. 'Assess' technologies require careful evaluation before committing.

Note: *These groups do not necessarily correlate with the CNCF maturity model (Sandbox, Incubating, and Graduated) which corresponds to the Innovators, Early Adopters, and Early Majority tiers from Geoffrey A. Moore's Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers.*

Sandbox: *Sandbox projects are in their earliest stages, meant for experimentation and foundational growth. They represent initial concepts and technology with significant room for evolution.*

Incubating: *Projects that have a solidified technical vision and growing contributor base but are still maturing in terms of community adoption, stability, and governance.*

Graduated: *Graduated projects are widely adopted and reliable. They have established a diverse community base supported by mature technical policies and governance.*

[1] *More information on these technologies can be found at the CNCF Landscape and LF AI & Data Landscape.*

TRENDS IN DEVOPS TECHNOLOGY USAGE

02

2. Multicluster Application Management Technologies

For multicluster application management (MCM) technologies, we see Argo CD and Cilium emerged as the two MCM technologies that respondents cumulatively would place in the “adopt” position of the technology radar.

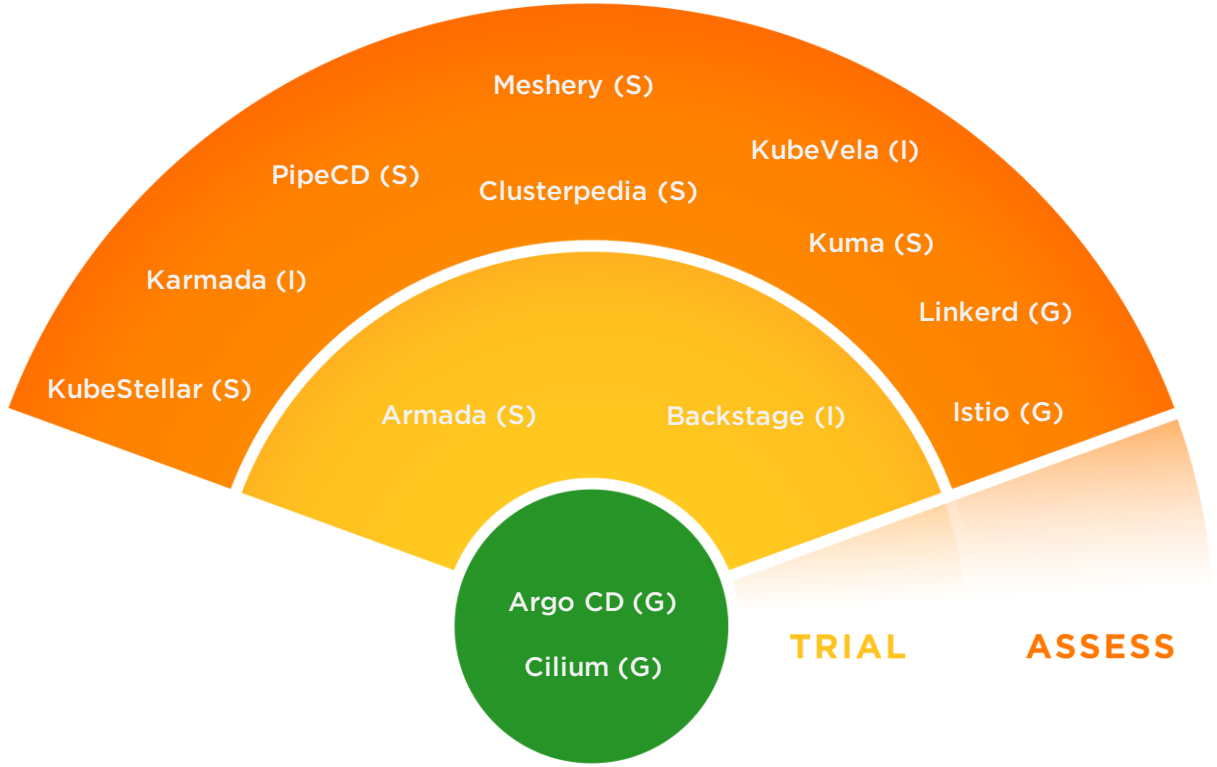

MCM Radar

G Graduated I Incubating S Sandbox

(n=204)

Based on developer perceptions: 'Adopt' technologies are considered reliable choices for most use cases, 'trial' technologies are worth exploring to see if they meet your specific needs and 'assess' technologies require careful evaluation before committing.

** Graduated, Incubating, and Sandbox refer to CNCF's hosted project levels*

2. Multicluster Application Management Technologies

Cilium received the highest usefulness score (+50), as well as receiving no negative, 1 or 2 star, ratings. Armada and Clusterpedia received very few negative ratings (2%, each), but a much smaller proportion of 5-star ratings (35% and 27%, respectively), suggesting that while developers don't consider them poorly matched to project requirements, they are not yet excelling, likely as a consequence of both being relatively new technologies in the sandbox stage.

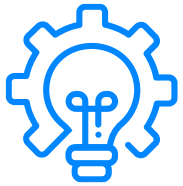

85%

THE PROPORTION OF 4 AND 5-STAR USEFULNESS RATINGS FOR ARGOCD

Usefulness ratings of multicluster application management technologies

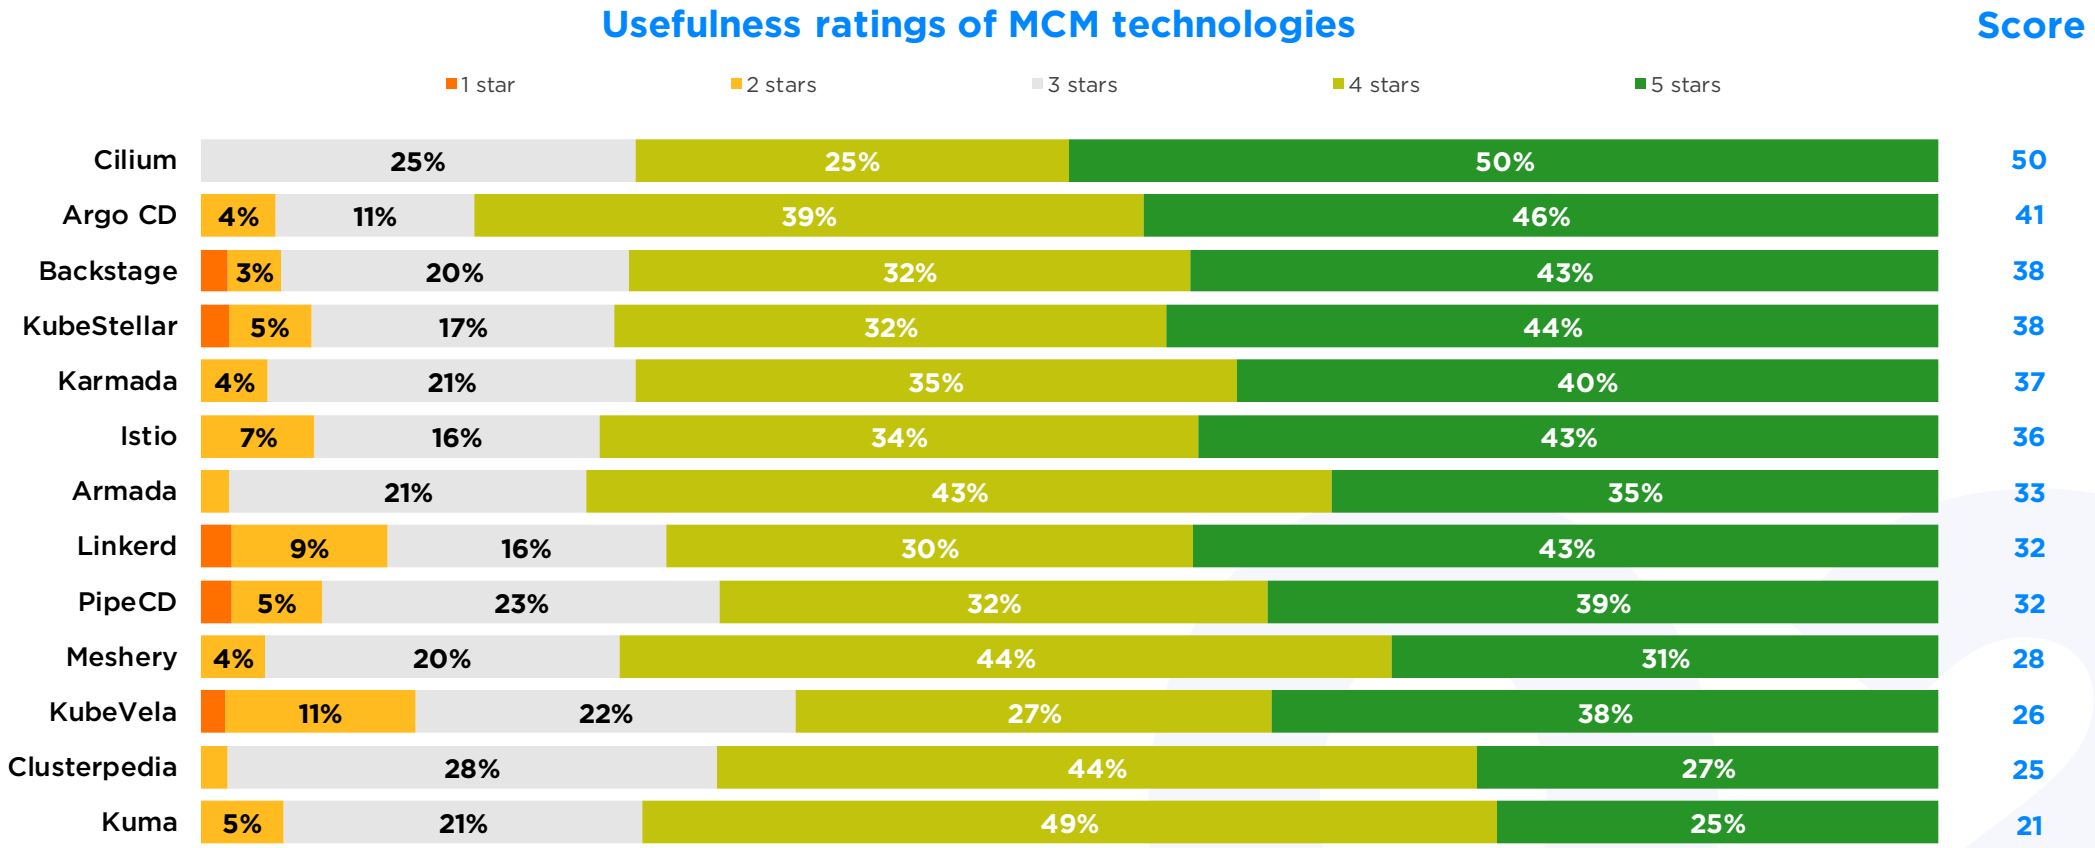

Question wording: How would you rate the following multicluster management tools/products with respect to these aspects (Usefulness)
% of developers familiar with each technology | Score (% of 5 star ratings minus the % of 1 and 2 star ratings) (n=204)

2. Multicluster Application Management Technologies

Cilium also received the highest maturity score (+47), cementing its position as a technology that the community considers both the most useful and mature. Despite its worse usefulness ratings and likelihood of recommendation, Linkerd receives high ratings for its maturity. As a project that has graduated from CNCF's project stages, this is a positive positioning. While not all technologies will meet the needs of all developer's projects, high levels of stability and reliability are key to ensuring a smooth baseline performance of a tool.

Half of developers asked gave Cilium a 5-star rating for maturity

Maturity ratings of multicluster application management technologies

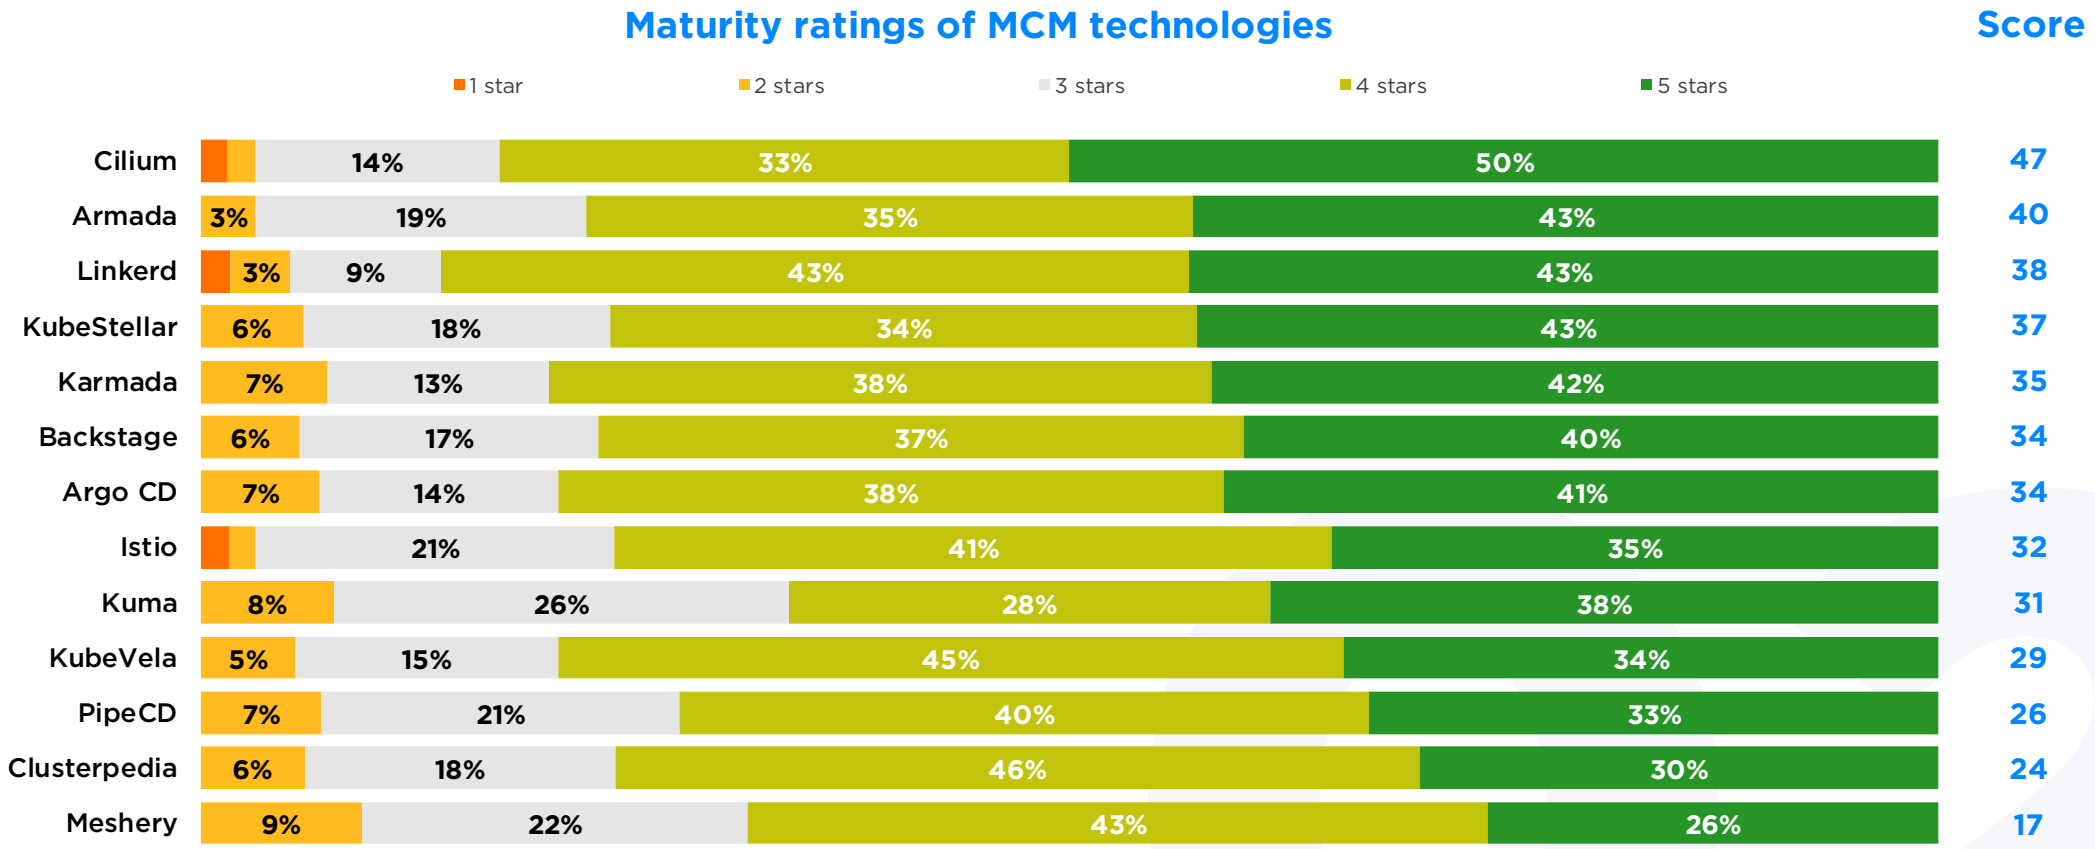

Question wording: How would you rate the following multicluster management tools/products with respect to these aspect (Maturity)
% of developers familiar with each technology | Score (% of 5 star ratings minus the % of 1 and 2 star ratings) (n=204)

2. Multicluster Application Management Technologies

Argo CD received the highest NPS, +87, which also corresponds to it being the most widely used MCM technology within our survey. Armada has the second-highest NPS, +79, which likely emerges from the high maturity ratings it received. During the personal usage of technologies, developers may be more likely to consider recommending a technology that is reliable, over one that can meet all the requirements of their projects. Downtime or failures may be viewed as a greater pain point than a tool overall being ideally matched to a project's requirements.

Despite this, Cilium receives a lower NPS (+70) than would be expected of its high usefulness and maturity ratings. This suggests there may be additional considerations that developers place on these technologies, before recommending them, that are not captured by our current research, and may benefit from further investigation.

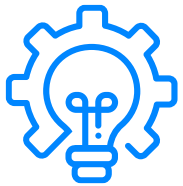

78%

OF CURRENT OR FORMER USERS OF MESHERY WOULD RECOMMEND IT TO OTHERS

Likelihood to recommend multicluster application management technologies

Likelihood to recommend MCM technologies

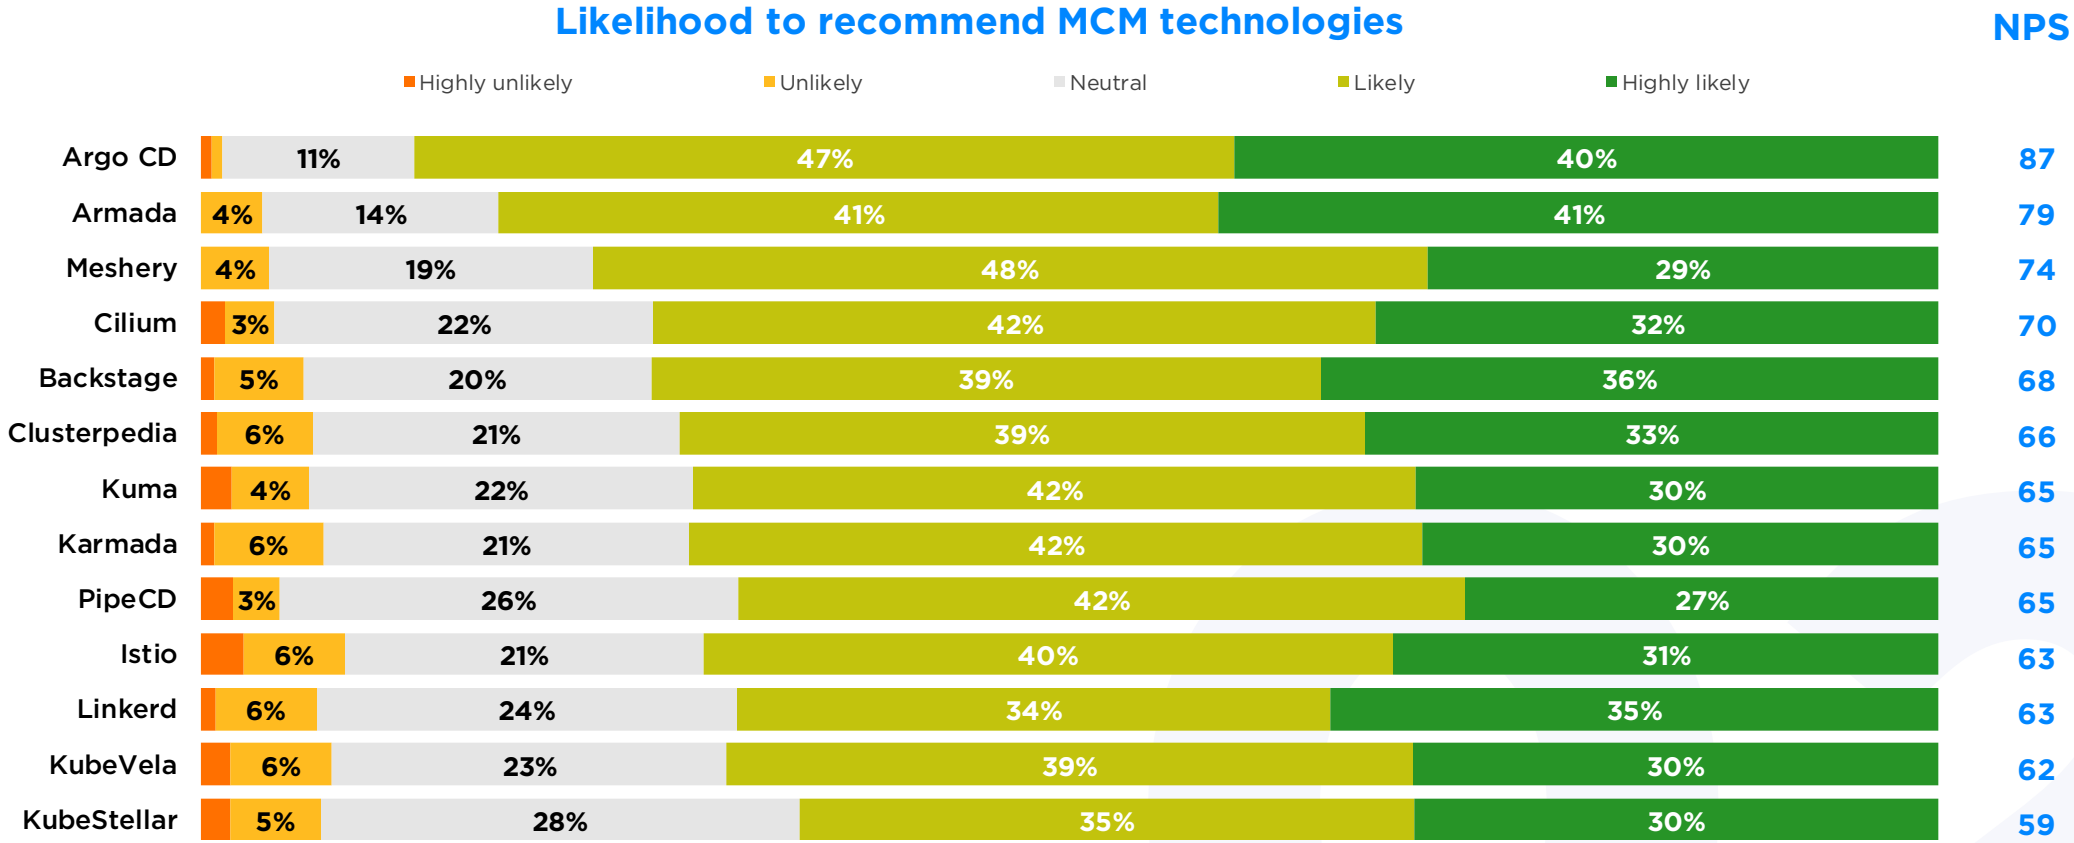

Question wording: How likely are you to recommend the following multicluster management technologies?
% of developers with experience with each technology | NPS (% of Likely and Very likely minus the % of Unlikely and Very unlikely) (n=204)

2. Multicluster Application Management Technologies

By using a normalised scale for each aspect of the technology considered for projecting their position on the technology radar, we can see the overall positioning of technologies by their radar position. In general, technologies in the adopt or trial position are more likely to receive higher normalised scores for each aspect, while those in the assess position receive lower scores.

However, it is possible to see technologies in the “assess” position performing well in some aspects, but in others, they performed worse, leading to their overall position. The technology radar positions are determined by their position across all factors, meaning that sporadic strong performance on one feature is not enough to move technology into the “adopt” position-banding.

Distribution of normalised scores for MCM technologies

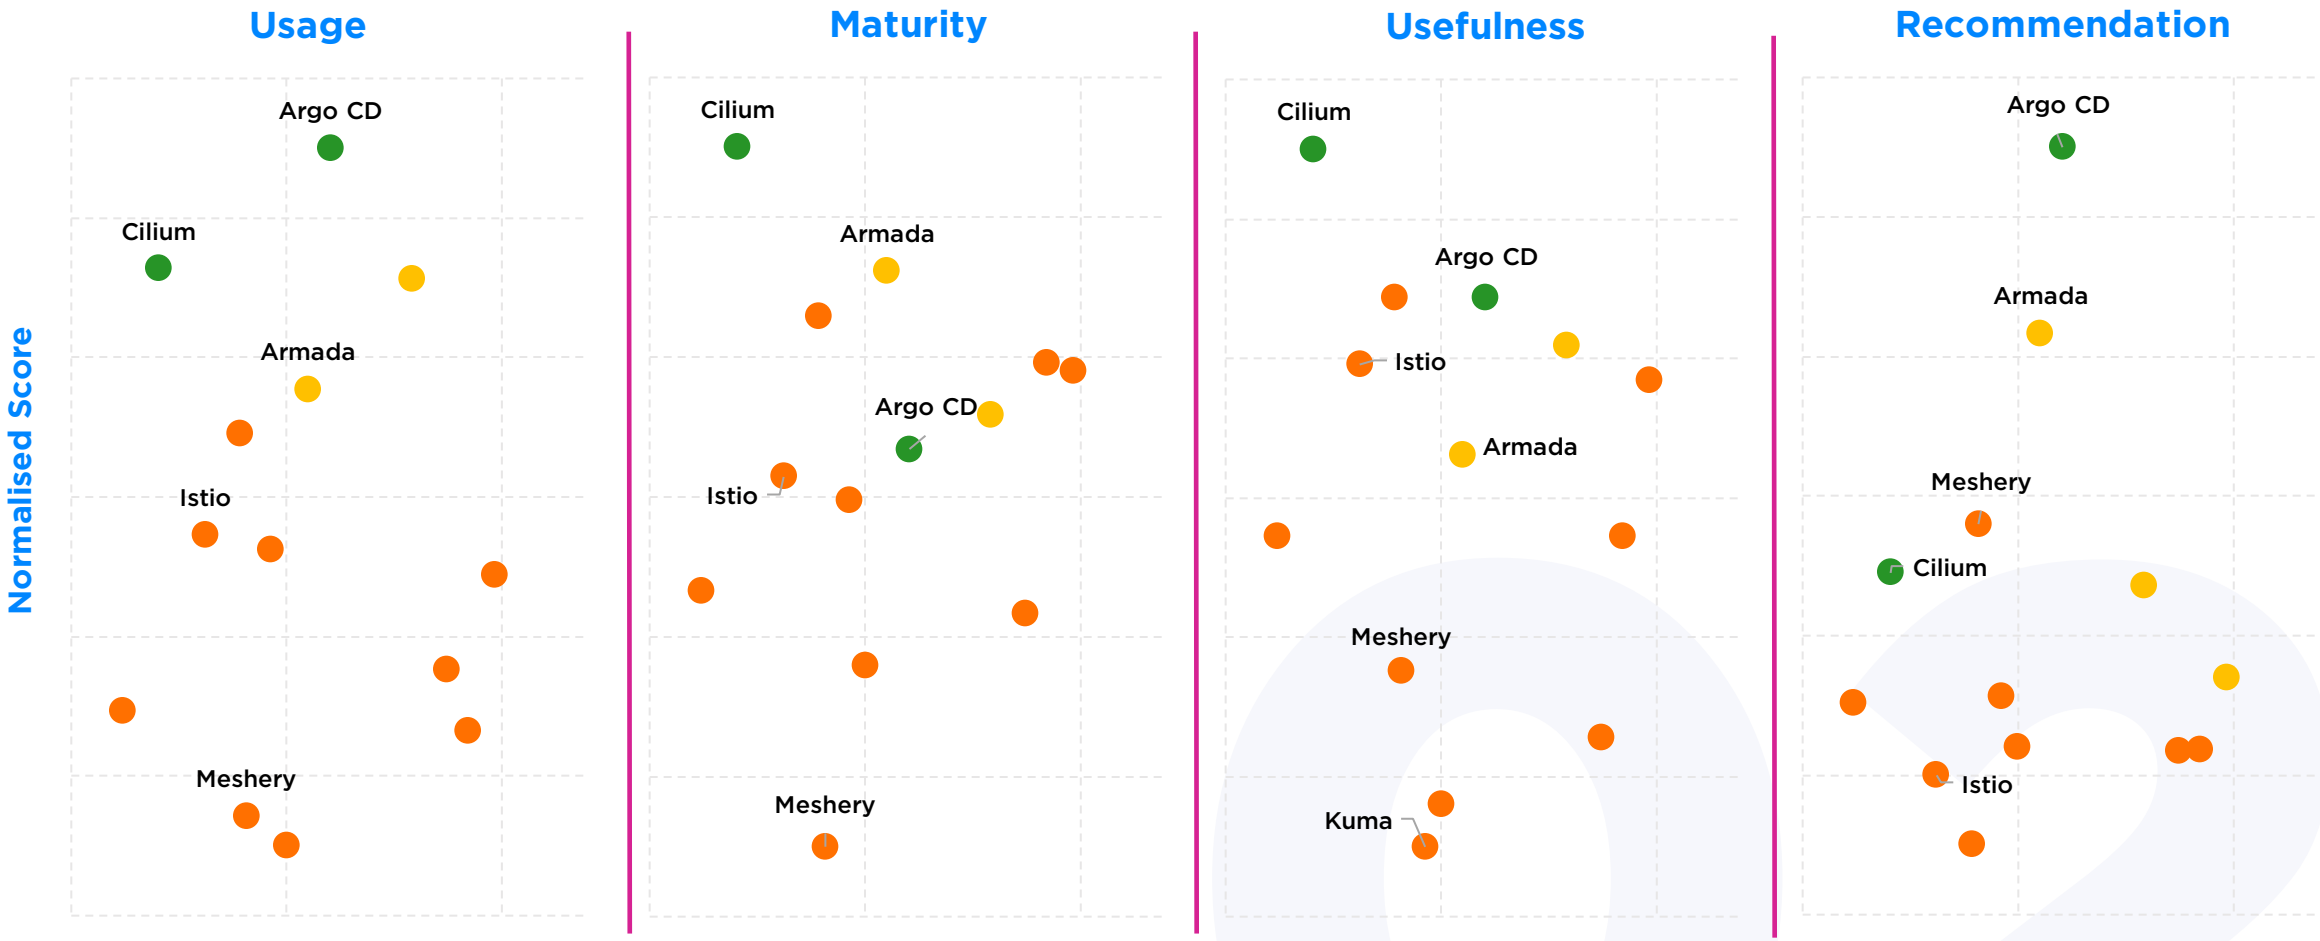

Legend: **Green:** Adopt, **Yellow:** Trial, **Orange:** Assess
Normalised score of usage, maturity, usefulness, and NPS (n=204)

BATCH/AI/ML COMPUTE TECHNOLOGIES

03

3. Batch/AI/ML Compute Technologies

Moving from multicluster application management technologies, we now look at batch/AI/ML compute technologies. Across all of the batch/AI/ML compute technologies developers were asked about, Apache Airflow was a clear standout. The most popular technology, with the highest maturity rating, and the second highest usefulness rating and NPS.

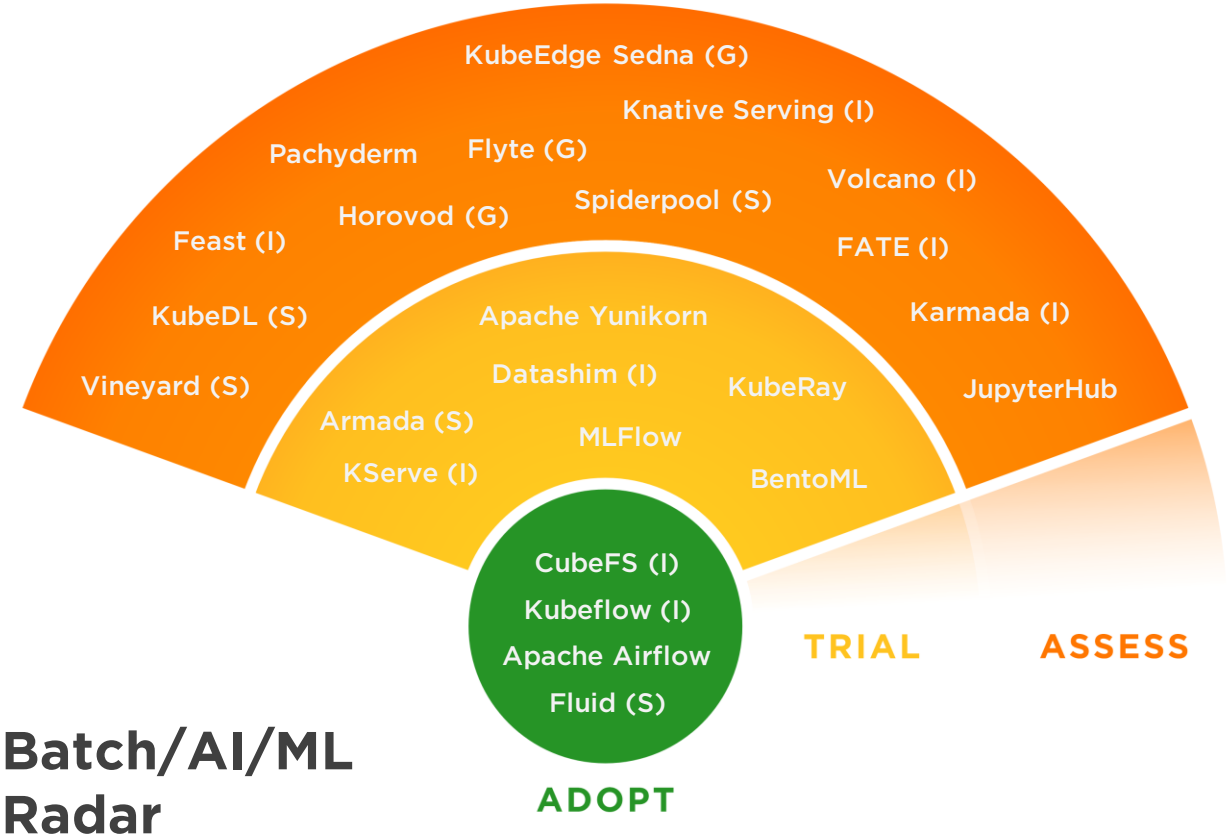

Batch/AI/ML Radar

G Graduated I Incubating S Sandbox

(n=217)

Based on developer perceptions: 'Adopt' technologies are considered reliable choices for most use cases, 'trial' technologies are worth exploring to see if they meet your specific needs and 'assess' technologies require careful evaluation before committing.

** Graduated, Incubating, and Sandbox refer to CNCF and LF AI & Data's hosted projects' levels*

3. Batch/AI/ML Compute Technologies

Kubeflow received the highest usefulness score (+47), but Apache Airflow (+40) received more four and five-star reviews. CubeFS (+40), Fluid (+36), and Kubeflow[1] achieved usefulness scores that placed them in the top five technologies, despite still being in either CNCF's sandbox or incubating stage of project maturity, indicating a promising future for these technologies. Spiderpool received a larger proportion of 5-star ratings (34%) than technologies with similar usefulness scores, such as JuptyrHub and Pachyderm, due to the higher proportion of 1 and 2-star ratings (20%). This suggests that for some developers Spiderpool is meeting their project requirements very well, but these may be more niche or less common use cases.

85% of MLFlow's ratings for usefulness were positive, but only 35% of them were 5-stars

[1] Kubeflow has multiple projects within it, such as ArgoCD and Kueue, that are largely invisible to the user but are a component of the developer experience of using it

Usefulness ratings of Batch/AI/ML technologies

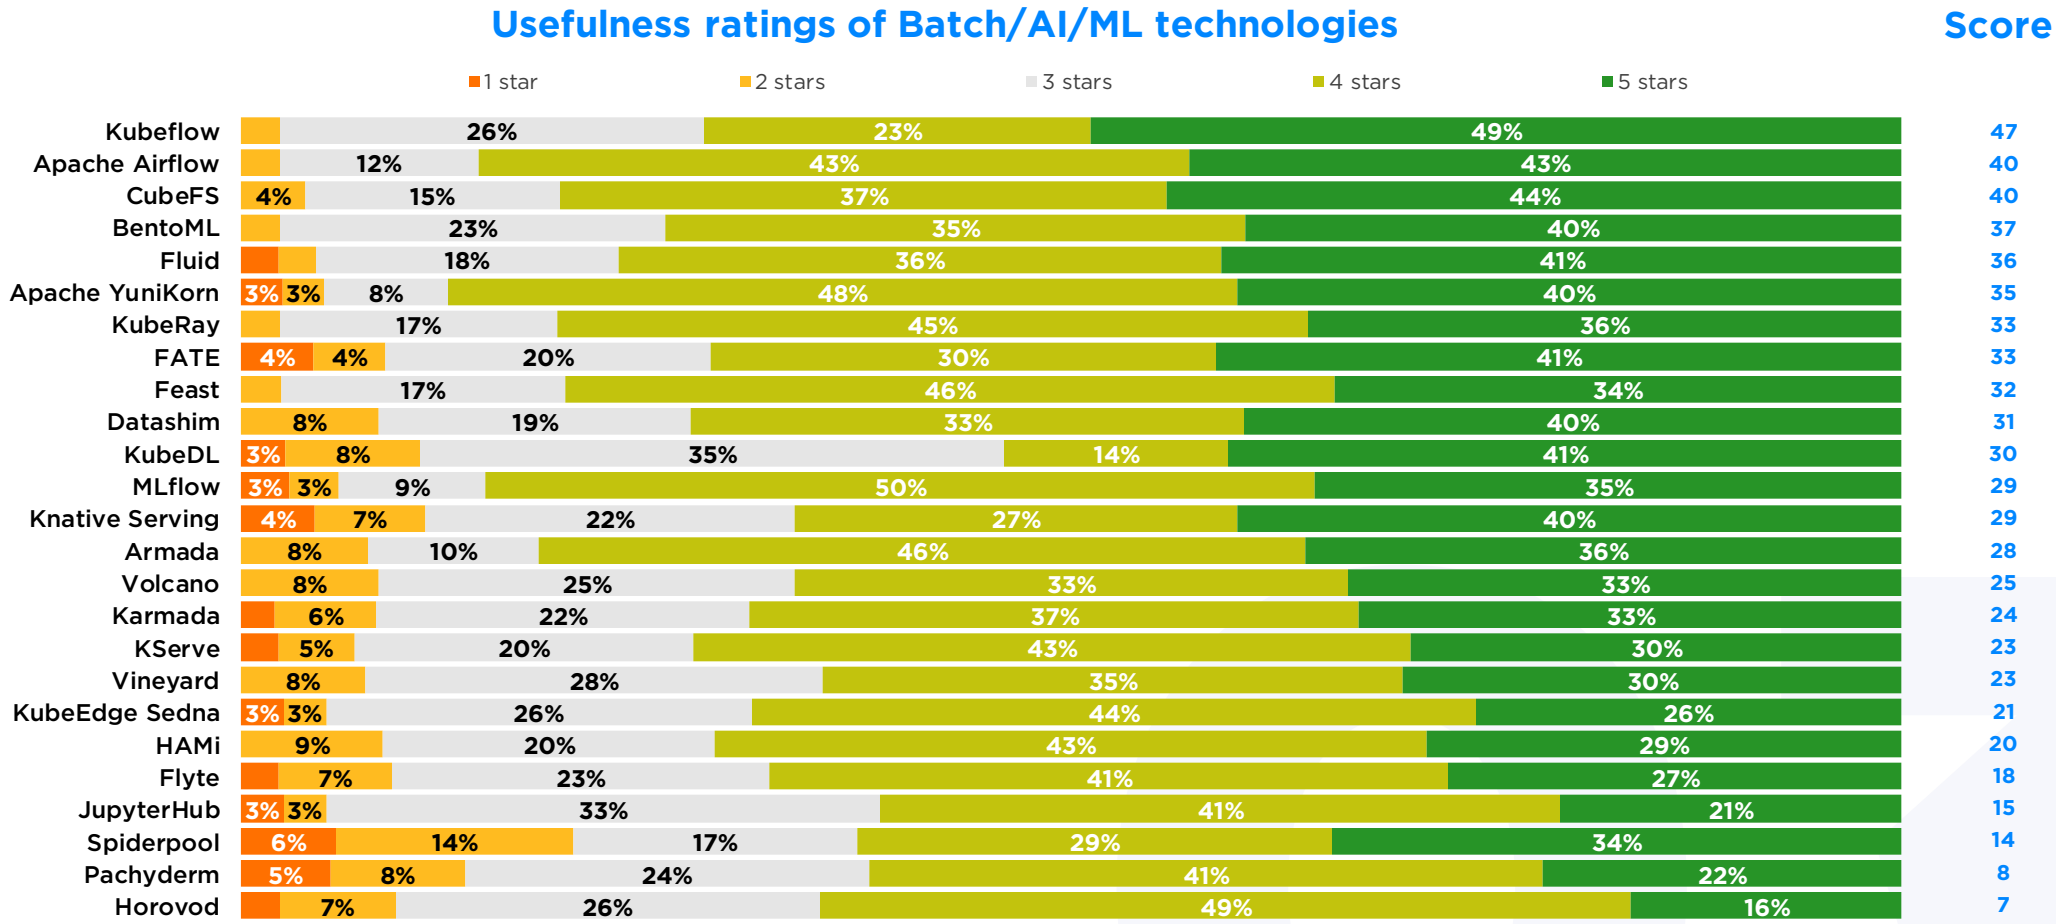

Question wording: How would you rate the following batch, AI, or ML computing products/tools with respect to these aspects? (Usefulness)
% of developers familiar with each technology | Score (% of 5 star ratings minus the % of 1 and 2 star ratings) (n=217)

3. Batch/AI/ML Compute Technologies

A staggering 88% of developers asked about Apache Airflow gave it either a 4 or 5-star rating for its maturity, with the majority of developers (58%) giving a 5-star rating. Only Apache Yunicorn receives a similar amount of overall positive ratings, but a much smaller percentage (32%) of 5-star ratings. KServe received a large percentage of negative ratings on its maturity (16%), indicating an important area where developers feel it is falling short.

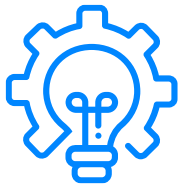

45%

OF DEVELOPERS FAMILIAR WITH CUBEFS GAVE IT A 5-STAR RATING ON MATURITY

Maturity ratings of Batch/AI/ML technologies

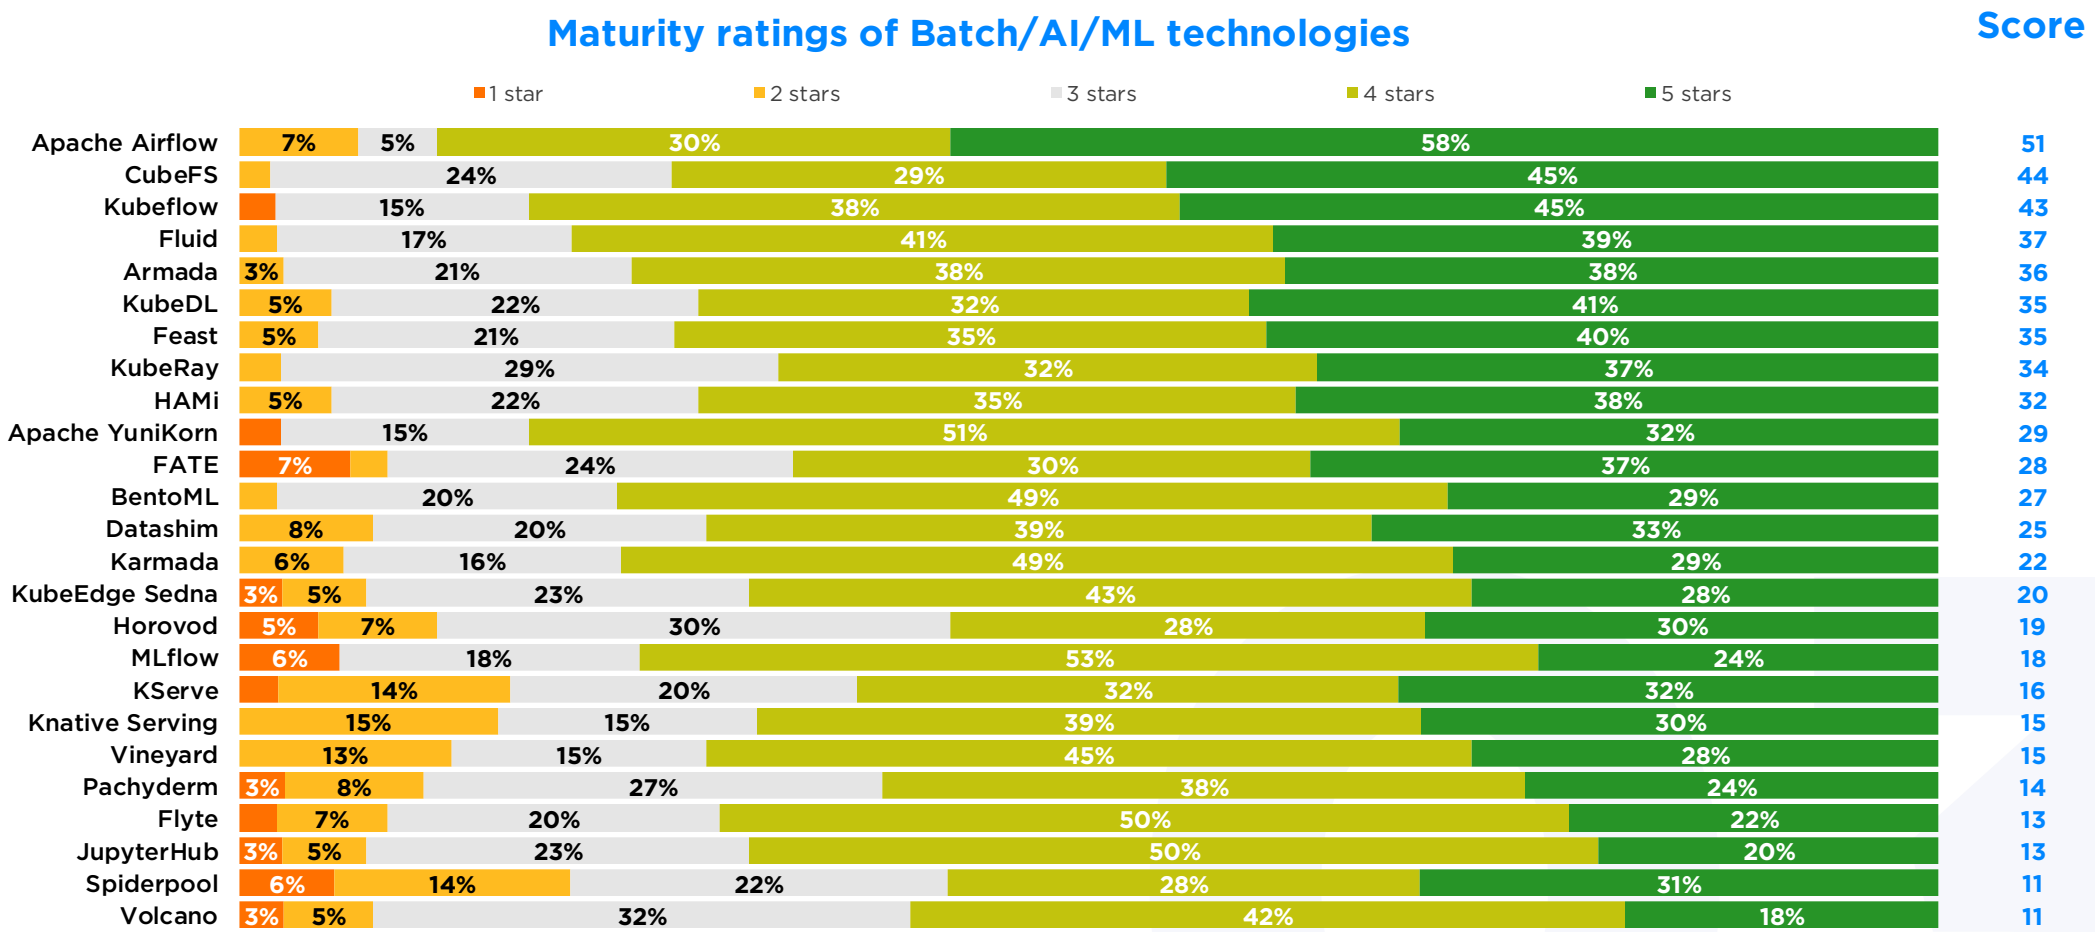

Question wording: How would you rate the following batch, AI, or ML computing products/tools with respect to these aspects? (Maturity)
% of developers familiar with each technology | Score (% of 5 star ratings minus the % of 1 and 2 star ratings) (n=217)

3. Batch/AI/ML Compute Technologies

Developers who indicated they had used a technology were asked how likely they were to recommend it to other developers. Despite the various ratings for maturity and usefulness, a clear majority of users of each technology recommend them. Vineyard received the fewest users who would recommend it, but this still constituted 66% of users surveyed. This, overall, makes recommendations a less effective metric in isolation for differentiating technologies. In general, developers provide favourable recommendations to the technologies they have invested time and energy in learning to use, reminding us of the importance of digging further into developer perceptions of technologies when constructing technology landscape radars.

Developers provide high overall recommendation ratings for tools they are familiar with. Vineyard received the fewest recommendations, but was still a majority among current or former users (66%)

Likelihood to recommend Batch/AI/ML technologies

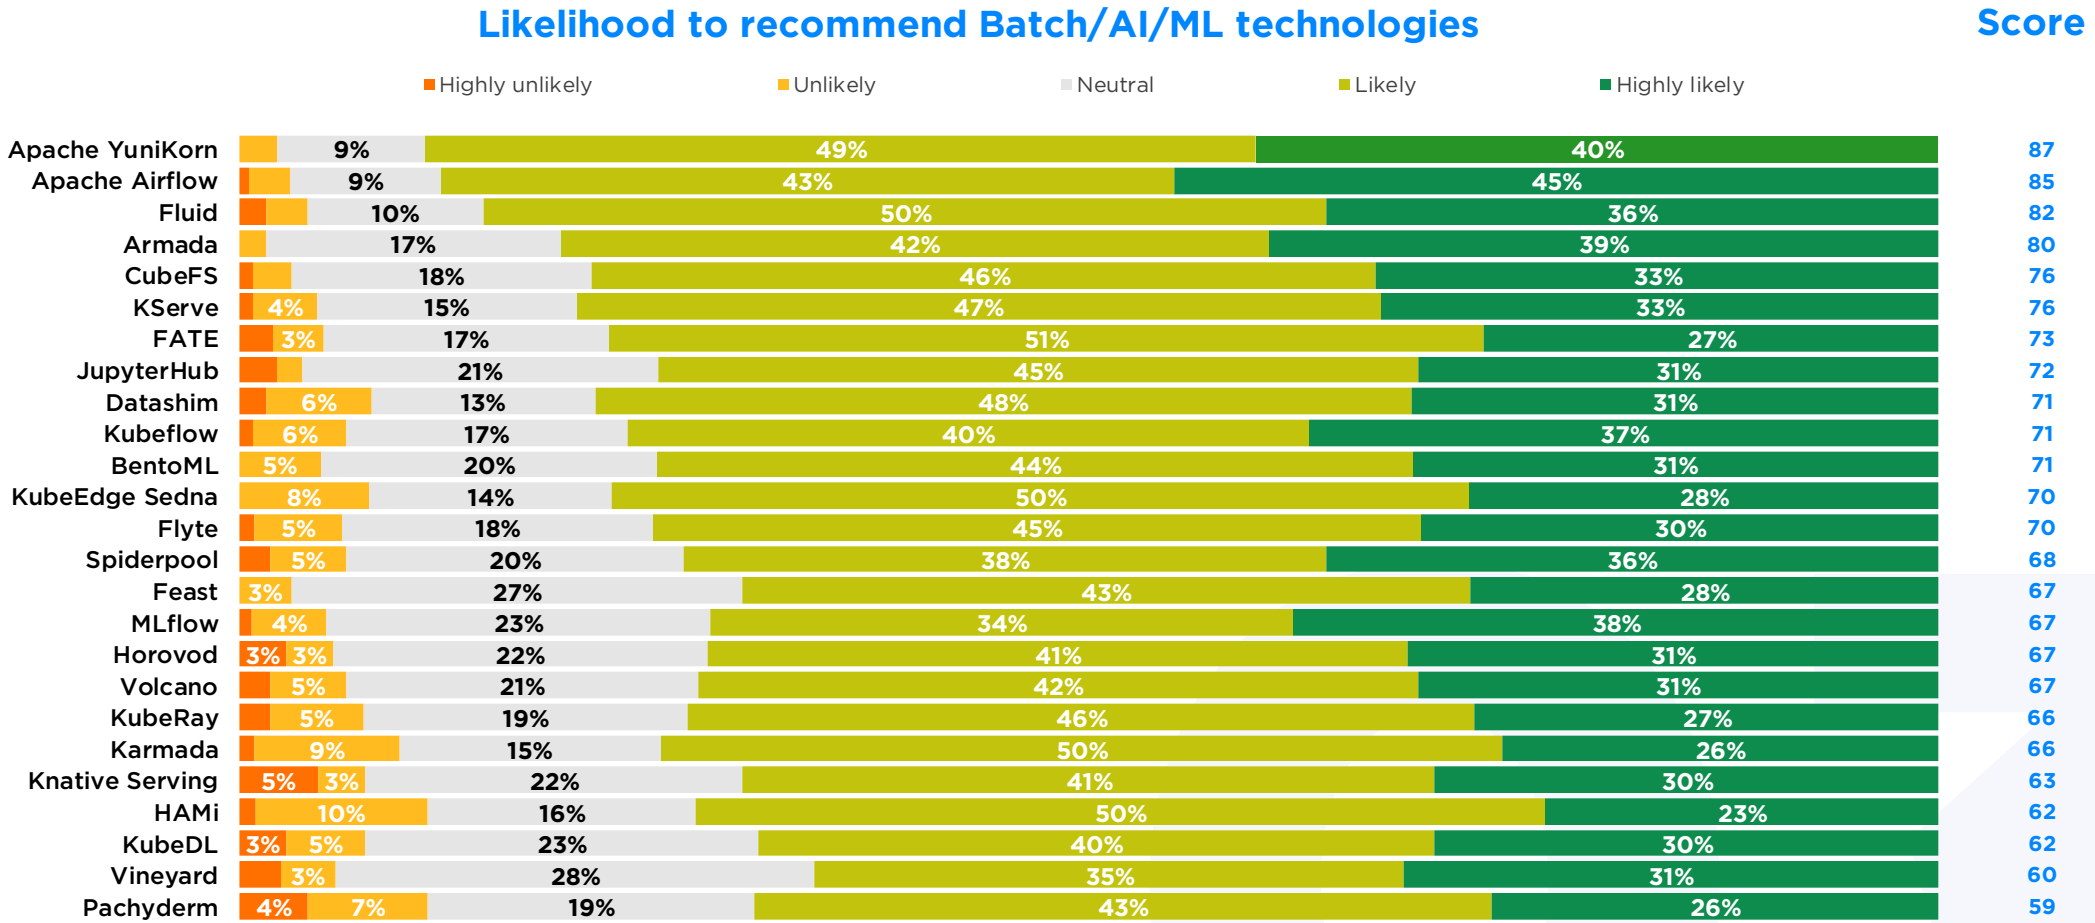

Question wording: How likely are you to recommend the following batch, AI, or ML computing technologies?
% of developers with experience with each technology | NPS (% of Likely and Very likely minus the % of Unlikely and Very unlikely) (n=217)

Distribution of normalised scores for Batch/AI/ML technologies

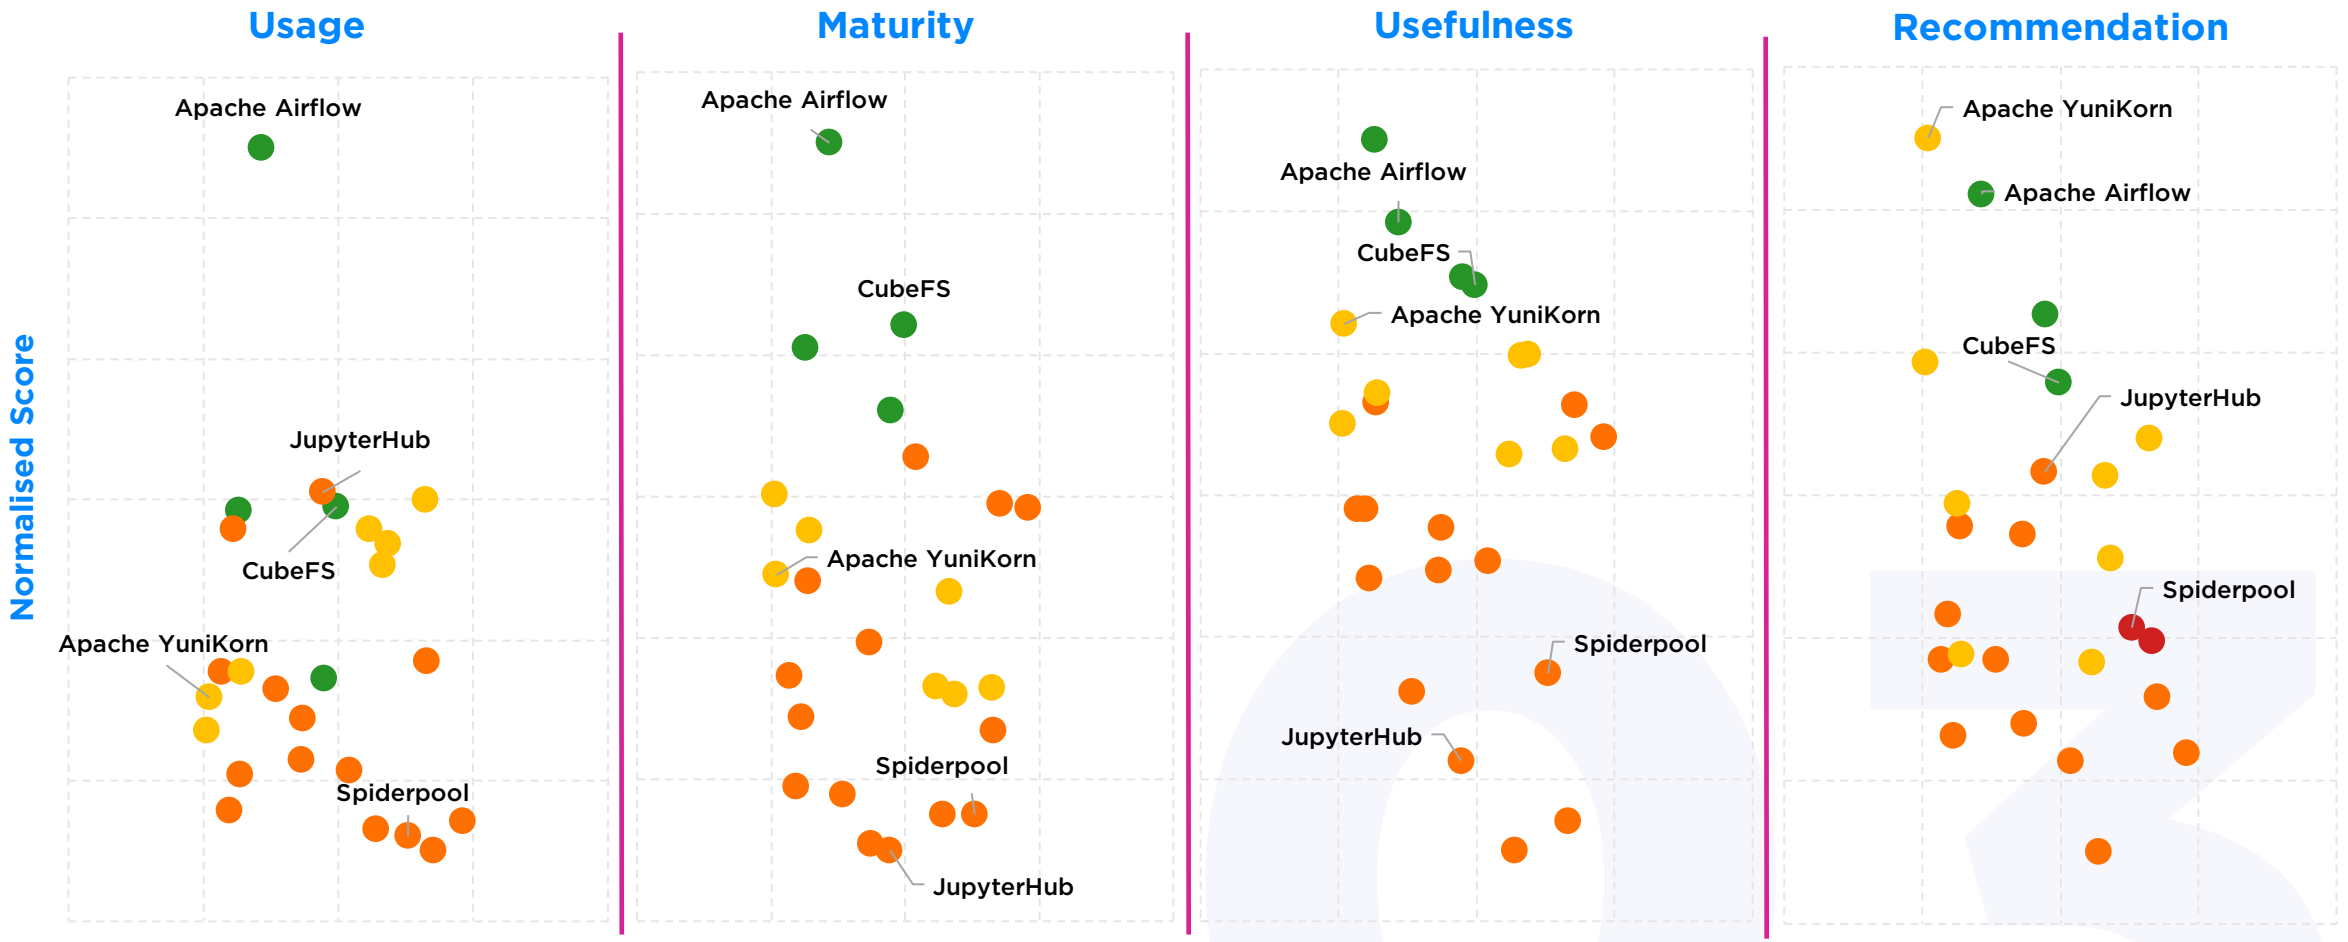

Legend: **Green:** Adopt, **Yellow:** Trial, **Orange:** Assess
Normalised score of usage, maturity, usefulness, and NPS (n=217)

METHODOLOGY

Methodology

Over two weeks between September and October 2024, more than 300 professional developers using cloud native technologies were asked for their opinions on various batch computing and multicluster application management technologies they were familiar with. Multicluster application management and batch/AI/ML compute technologies were identified by CNCF as two technology areas to investigate further for this technology landscape radar. The individual technologies selected are based on those identified by CNCF's End User Technical Advisory Board (TAB).

From the developer responses four variables were derived:

- **Weighted usage score:** Using the proportion of developers familiar with a technology, longer-term usage was weighted higher than more recent adoption, indicating long-term viability and continued use of a technology
- **Maturity score:** Developers were asked to rate technologies on a 5-star scale of maturity, where 1 star indicated low maturity and 5 stars indicated high maturity. The maturity score for each technology was calculated as the percentage of 5-star ratings minus the percentage of 1- and 2-star ratings, multiplied by 100 to create a scale of -100 to 100.

Methodology

- **Usefulness score:** Developers were asked to rate technologies on a 5-star scale of usefulness, where 1 star indicated low usefulness and 5 stars indicated high usefulness. The usefulness score for each technology was calculated as the percentage of 5-star ratings minus the percentage of 1- and 2-star ratings, multiplied by 100 to create a scale of -100 to 100.
- **Net Promoter Score (NPS):** Developers were asked to state how likely they were to recommend the technology, on a scale of highly likely, likely, neutral, unlikely, highly unlikely. This was converted to an NPS by taking the percentage of likely and highly likely and subtracting the percentage of unlikely and highly unlikely. This was multiplied by 100 to create a scale of -100 to 100.

All three variables were normalised to a scale of 0 to 1, after which multidimensional scaling (MDS) was used to visualise the levels of similarity and difference between each of the technologies. MDS was chosen to provide a proximity-based analysis to identify similarly performing technologies. The results of this MDS analysis was then clustered into four distinct groups to provide the adopt, trial and assess bands of the technology radar.

Subjective nature of Likert scales

In our research, we employed Likert scales to capture developers' opinions on the maturity and usefulness, from 1 to 5-stars, of the various multicluster application management and batch computing technologies surveyed. While these ratings are inherently subjective, reflecting individual perceptions and experiences, they provide valuable insights into the developer community's views. The nature of our research is centred on investigating developer perceptions of these aspects, making the subjective nature of the ratings not only acceptable but also valuable for our analysis. Although the subjective nature of Likert scales may influence the interpretation of results, as different respondents may have varying standards for rating, this variability enriches our understanding of the developer experience.

Despite these nuances, analysing the distribution of ratings — such as the difference between the number of 5-star ratings and those of 1 and 2 stars — serves as a practical measure for understanding developer sentiments. This approach allows us to identify trends and patterns that can inform decision-making, highlighting areas of strength and opportunities for improvement within the surveyed technologies. Thus, we assert that Likert scales are an effective tool for gauging developer perceptions and experiences.

Respondent demographics

Respondents were initially asked about where their projects ran or were deployed, to identify their position as a 'cloud developer'. Following this, they were asked which technologies they were currently using, that we associate with cloud-native development approaches, including technologies such as Infrastructure as Code, service meshes, and serverless computing.

Respondents were recruited from a combination of third-party panels, and within CNCF's community. For privacy and data minimisation purposes, exclusion is based on internal consistency and survey-taking behaviour metrics. As such, information on the organisation the respondent works for is not carried through to any analysis. This privacy also helps encourage greater honesty from respondents, who do not have concerns that their expressed opinion will be associated with them.

Due to the nature of third-party panels making up the significant majority of respondents, we consider the risk of multiple respondents from the same organisation responding to be low, and as such do not engage in deduping cleanses. However, should more than one individual from the same organisation respond to the survey, we do not consider it to impact the validity of the results.

Within the same organisation, developers may be using different technologies. Further, while usage was used in the determination of each technology's position on the technology landscape radar, the developer's personal perceptions corresponded to 75% of the score the technology received.

Methodology

Size of organisation respondents are working for

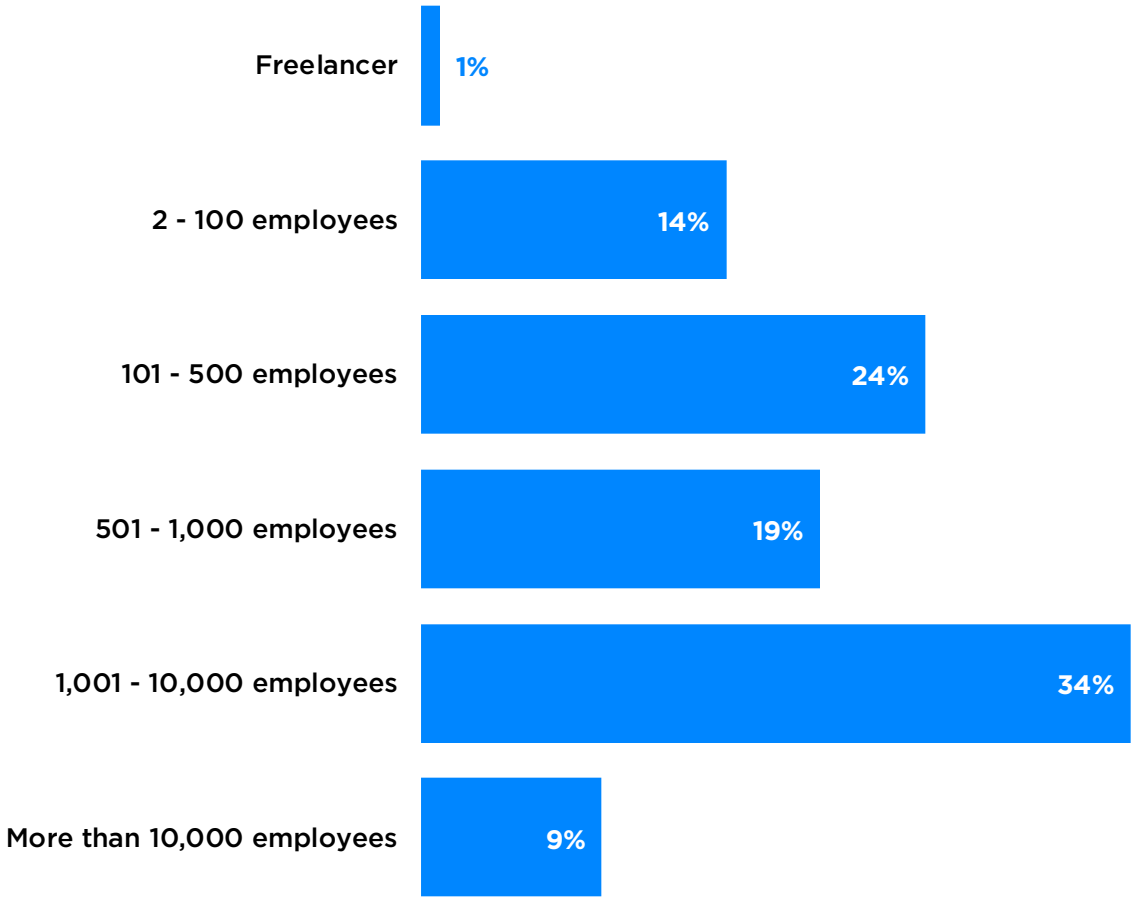

% of developers | % of developers (n=340)

Areas of development respondents are working within

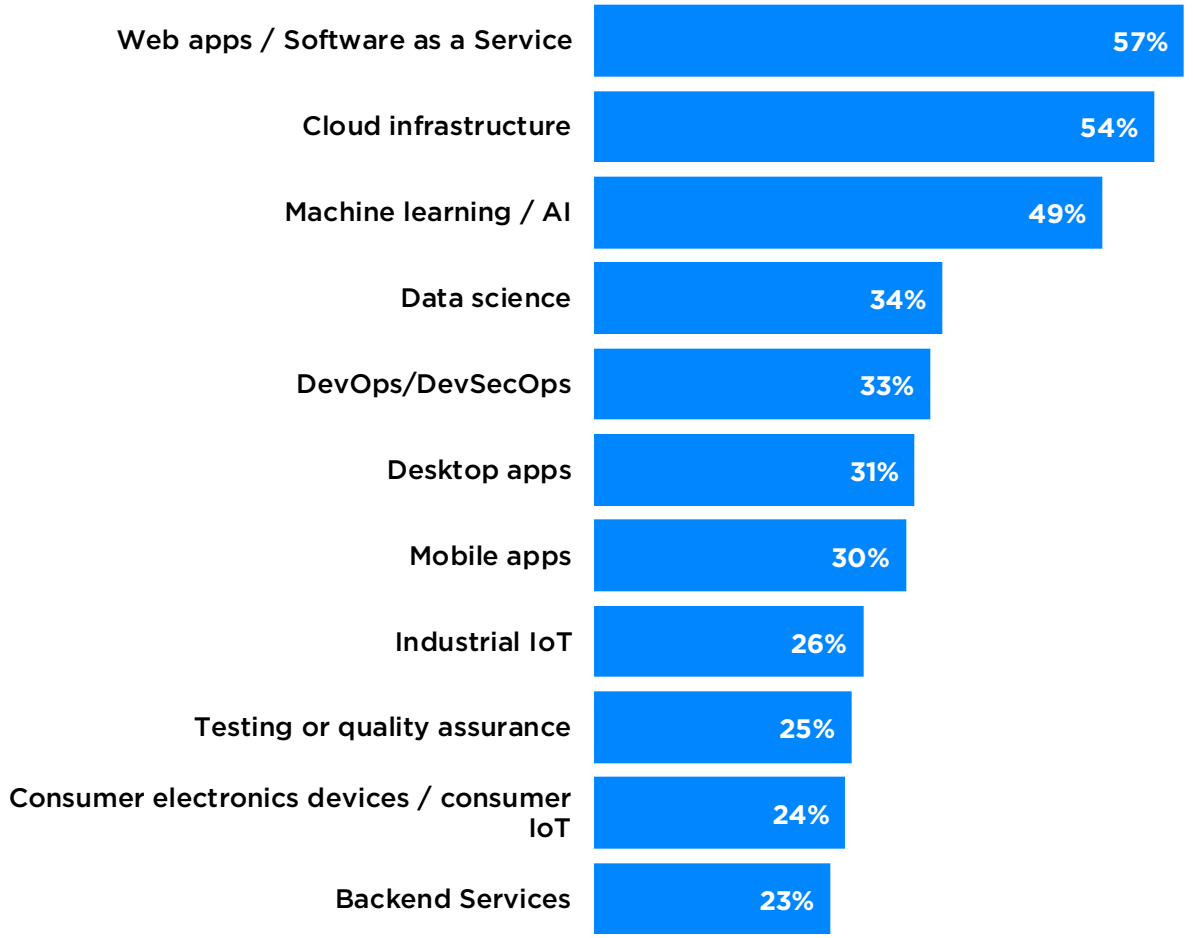

Methodology

Industry verticals developers are involved in

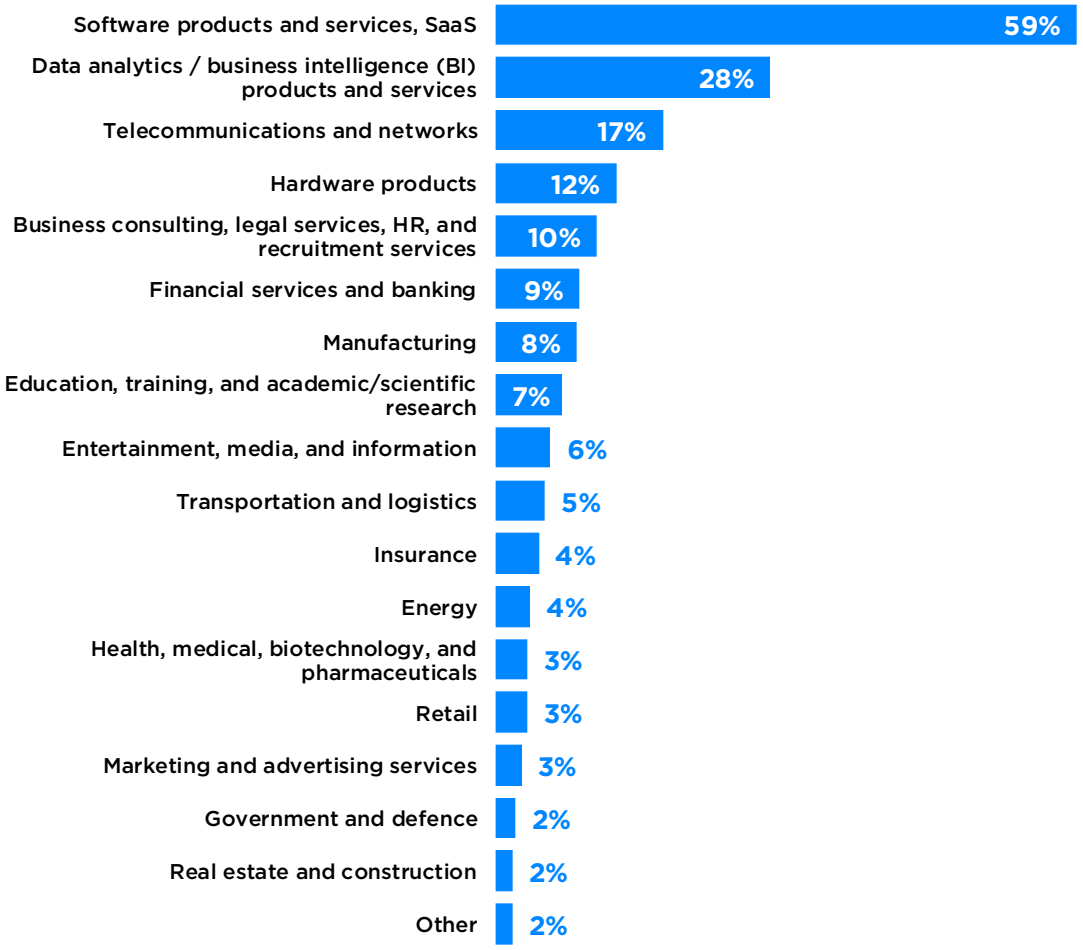

% of developers | % of developers (n=340)

Developer self-assessment of perception to new cloud-native technologies

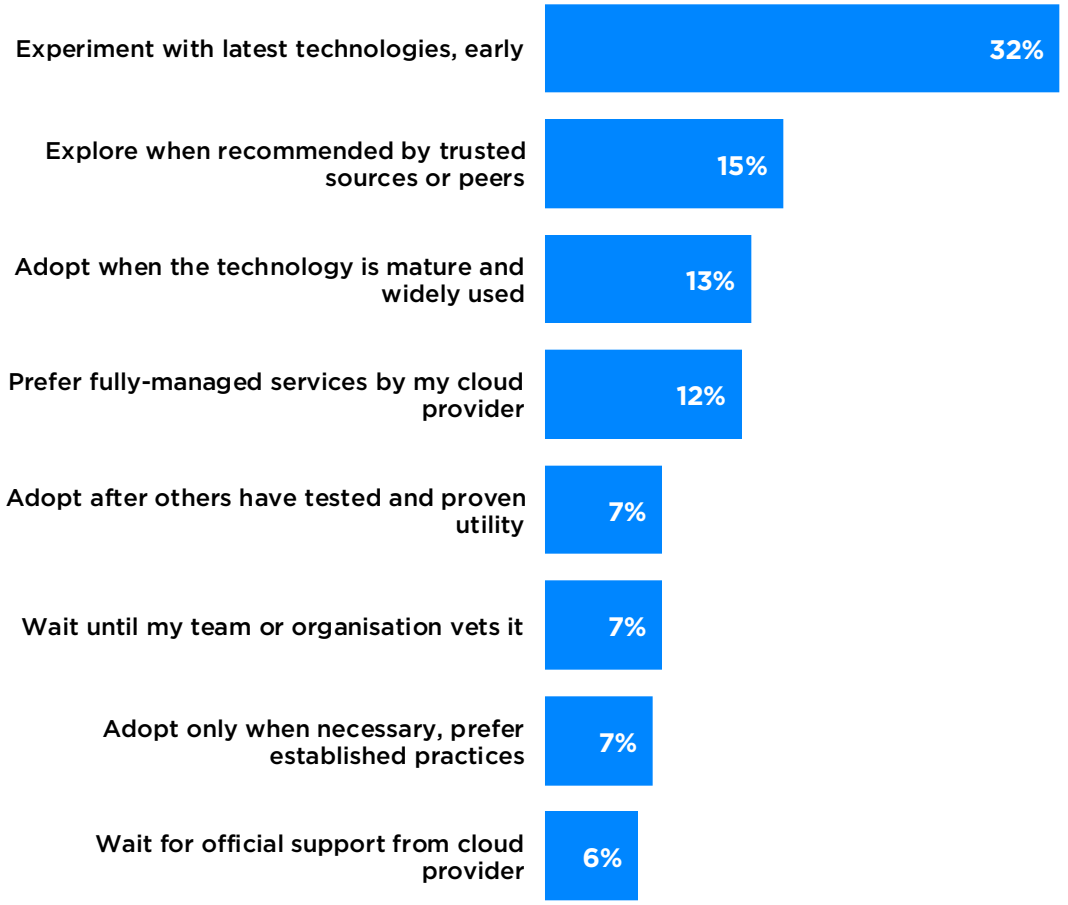

About SlashData

/DATA

Understand developers. Inspire the future of technology.

We survey 30,000+ developers annually – across Web, Desktop, Cloud, Mobile, Industrial IoT, AR/VR, Machine Learning and Data Science, Games, Consumer Electronics and Apps/Extensions for 3rd party ecosystems - to help companies understand who developers are, what they buy and where they are going next.

WHO DEVELOPERS ARE

Developer population sizing
Developer segmentation

WHAT THEY BUY

Why developers are adopting competitor products – and how you can fix that

WHERE THEY ARE GOING

Emerging platforms – augmented & virtual reality, machine learning

https://www.slashdata.co


/DATA
CLOUD NATIVE
COMPUTING FOUNDATION