2021 STATE OF MULTICLOUD

A CIO's Guide to the Underlying Dynamics Fueling Multicloud Strategies

PREFACE

The origins of this survey began with a single question: If application workloads could someday move freely across clouds, what would be the primary driver? And then, of course, we had to ask, what are the barriers to that world becoming a reality?

This survey aims to capture the state of multicloud. The data covers where we are today, of course. But it also aims to understand where we collectively want to be. Technology choices are never made without tradeoffs. Understanding broader aspirations gives us the context to better understand *why* people make the choices they make.

Will we ever realize a world where workloads freely move across clouds? We'll let you debate that with your friends. But, where would we be if we made technology choices without a little imagination? Not here, that's for sure.

Constraints are a reality. But resigning ourselves to the status quo need not be. Pushing through barriers, however incrementally, and solving the hard problems—that requires imagination from all of us.

Every year—every day, really—you are making hard choices as you apply yourself and your teams to building the applications, services, and platforms that are transforming the business. It's a journey, it's a challenging one, and everyone is on it. We hope the findings are both interesting and valuable to you as you progress on this journey.

Naturally, some are further along than others, which is why our analysis is at times done with a particular lens on leaders, those on par with the majority, and laggards. We can all learn from each other.

On that note, we would like to wholeheartedly thank the survey respondents. Without your time and candor this would not be possible.

Thank you.

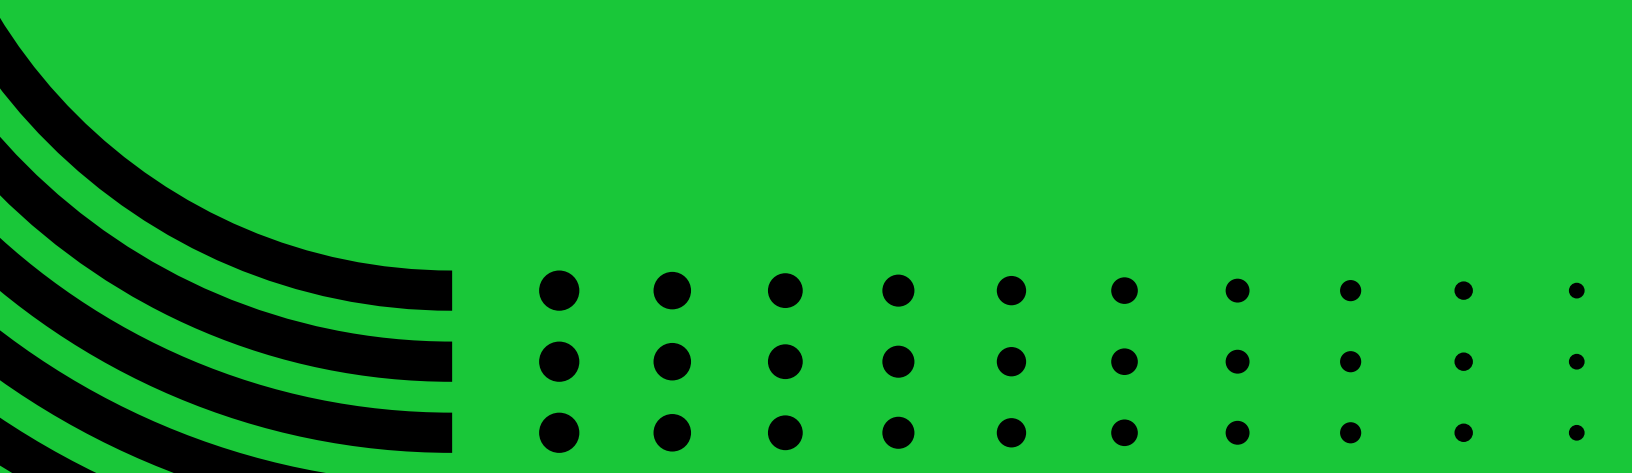

Table of Contents

Methodology & Demographics

We conducted our survey in January of 2021 and garnered responses from 819 participants. Of these, 36% were in an Infrastructure/Operations function, 19% were in an Architecture function, and the remaining 47% were a mix of Line-of-Business/Non-IT (9%), Applications (11%), Development (8%), DevOps (6%), Container Platforms (3%), and Security/Compliance (9%). The participants represent the full range of company sizes: Over 5000 (40%), 1,000-4,999 (22%), 200-999 (19%), 1-199 (18%), and Self-employed (1%).

In terms of industry, respondents were in Technology (21%), Service Provider (10%), Healthcare (12%), Financial Services (13%), Government (7%), Education (9%), Retail/E-Commerce (8%), Media & Entertainment (3%), and Other (17%). Role distribution was primarily an even mix between Individual Contributors (45%) and Director/Managers (45%); Executives constituted 11%. Lastly, for gender, 85% of respondents were male, 9% female, 1% non-binary, and 6% preferred not to say.

Highlights

1. Multicloud leaders aspire to leverage the right apps and services to differentiate their business, but **in 2021 they value business leverage** over the benefits of any one cloud.
2. In 2021 **leaders prioritize launching new internal services**, indicating progress on the digital transformation journey.
3. **For 61% of organizations** containerization will play a strategic role within 18 months, today it is already strategic for nearly 20%.
4. **56% are using at least one container platform today,** 79% of which are using commercial Kubernetes.
5. 56% of those on their container/cloud native journey are running containers in production, but **complexity at scale is hindering advanced use cases.**
6. **Optimization (25%) is the most important initiative** for organizations adopting public cloud in the coming year, followed by advancing a multicloud strategy (21%).
7. 30% of organizations are using **3 or more clouds** today.
8. **83% of organizations** are using at least one cloud managed service today
9. **For 62% of organizations** public cloud PaaS will play a strategic role for their business within 18 months.
10. **Complexity at 38% is the leading barrier** by far to edge computing becoming conventional.

What is multicloud?

We would love to offer up a precise definition of multicloud. But that would be about as helpful as pointing out that using the term "on-**premise**" in reference to grounds or buildings is an incorrect use of the English language. Regardless of whether you believe multicloud is when workloads move dynamically across clouds, or that it simply means having a footprint in multiple clouds, or something in between, we believe that there is a perspective for you in this analysis.

Our Lens on Leaders

Throughout this report we analyze the data by comparing "leaders" versus those "on par with the majority" versus "laggards." We asked respondents to self-identify based on the question: *"When it comes to leveraging new technologies to advance business goals, my organization is:"* Thirty four percent (34%) considered their organization a leader, 53% considered themselves on par with the majority, and 13% were self-described laggards.

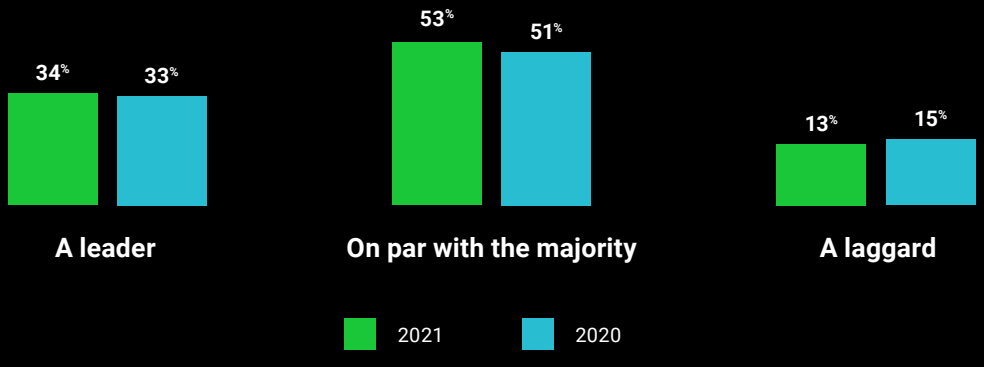

CHAPTER 1:

MULTICLOUD

Multicloud leaders aspire to leverage the right apps and services to differentiate their business.

There are many considerations that shape why and how organizations implement multicloud. Since 2019, however, leaders have consistently cited the ability to leverage different application services as the primary driver of application workloads "moving freely across clouds." 2021 is no different. As we have stated from the onset of this annual report, when clouds compete, customers win.

While multicloud is still somewhat nascent—only 30% of organizations are using three or more clouds (**see fig. 35**) —in the coming year advancing a multicloud strategy is the most important initiative for 21% of organizations overall, second only to optimizing existing cloud resources at 25%. For leaders the numbers are 28% and 25% respectively (**see fig. 33**).

We're already seeing cloud providers having to compete on their services. Customers evaluate cloud services based on what's best for their applications and their business. But, in some cases, cloud providers with a "best-of-breed service" will offer volume discounts on their increasingly commoditized IaaS. Even Kubernetes-as-a-Service (KaaS) offerings are falling into that "commoditized infrastructure" definition. Every cloud has its K8s flavor, but the native portability of this container platform makes it much easier to migrate a modern containerized application from one provider to another.

We are by no means at the point where workloads are dynamically moving across clouds (data gravity is the next big problem to solve, **see fig. 7**), but there is enough flexibility to maintain leverage.

Question: Do you believe that one day application workloads will move freely across clouds? If so, what will be the primary driver? (N = 819)

Figure 1 – *Overall: 2019-2021*

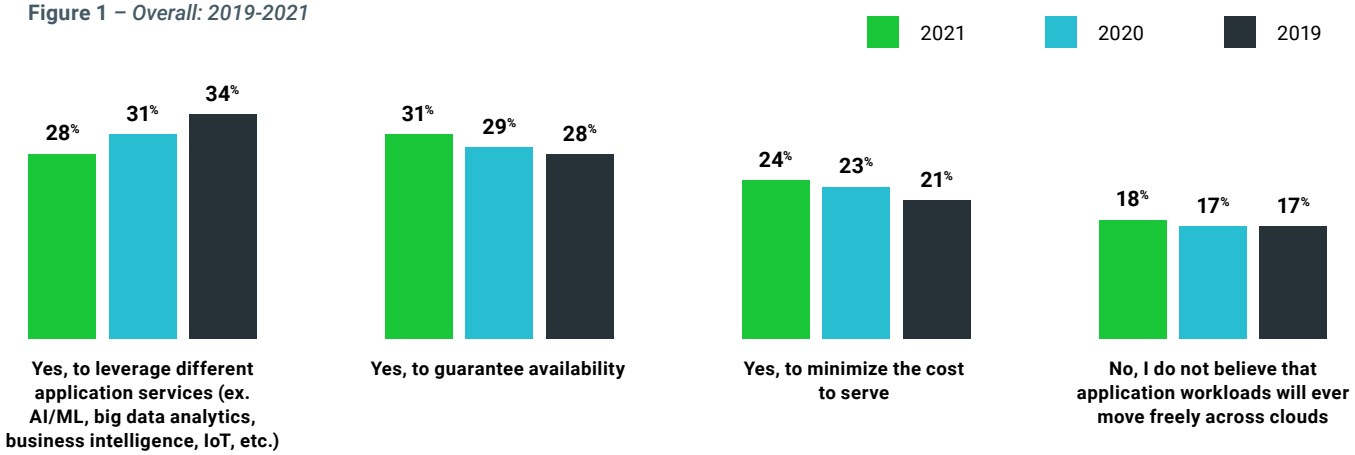

Figure 2 – *Leaders vs. Laggards: 2021*

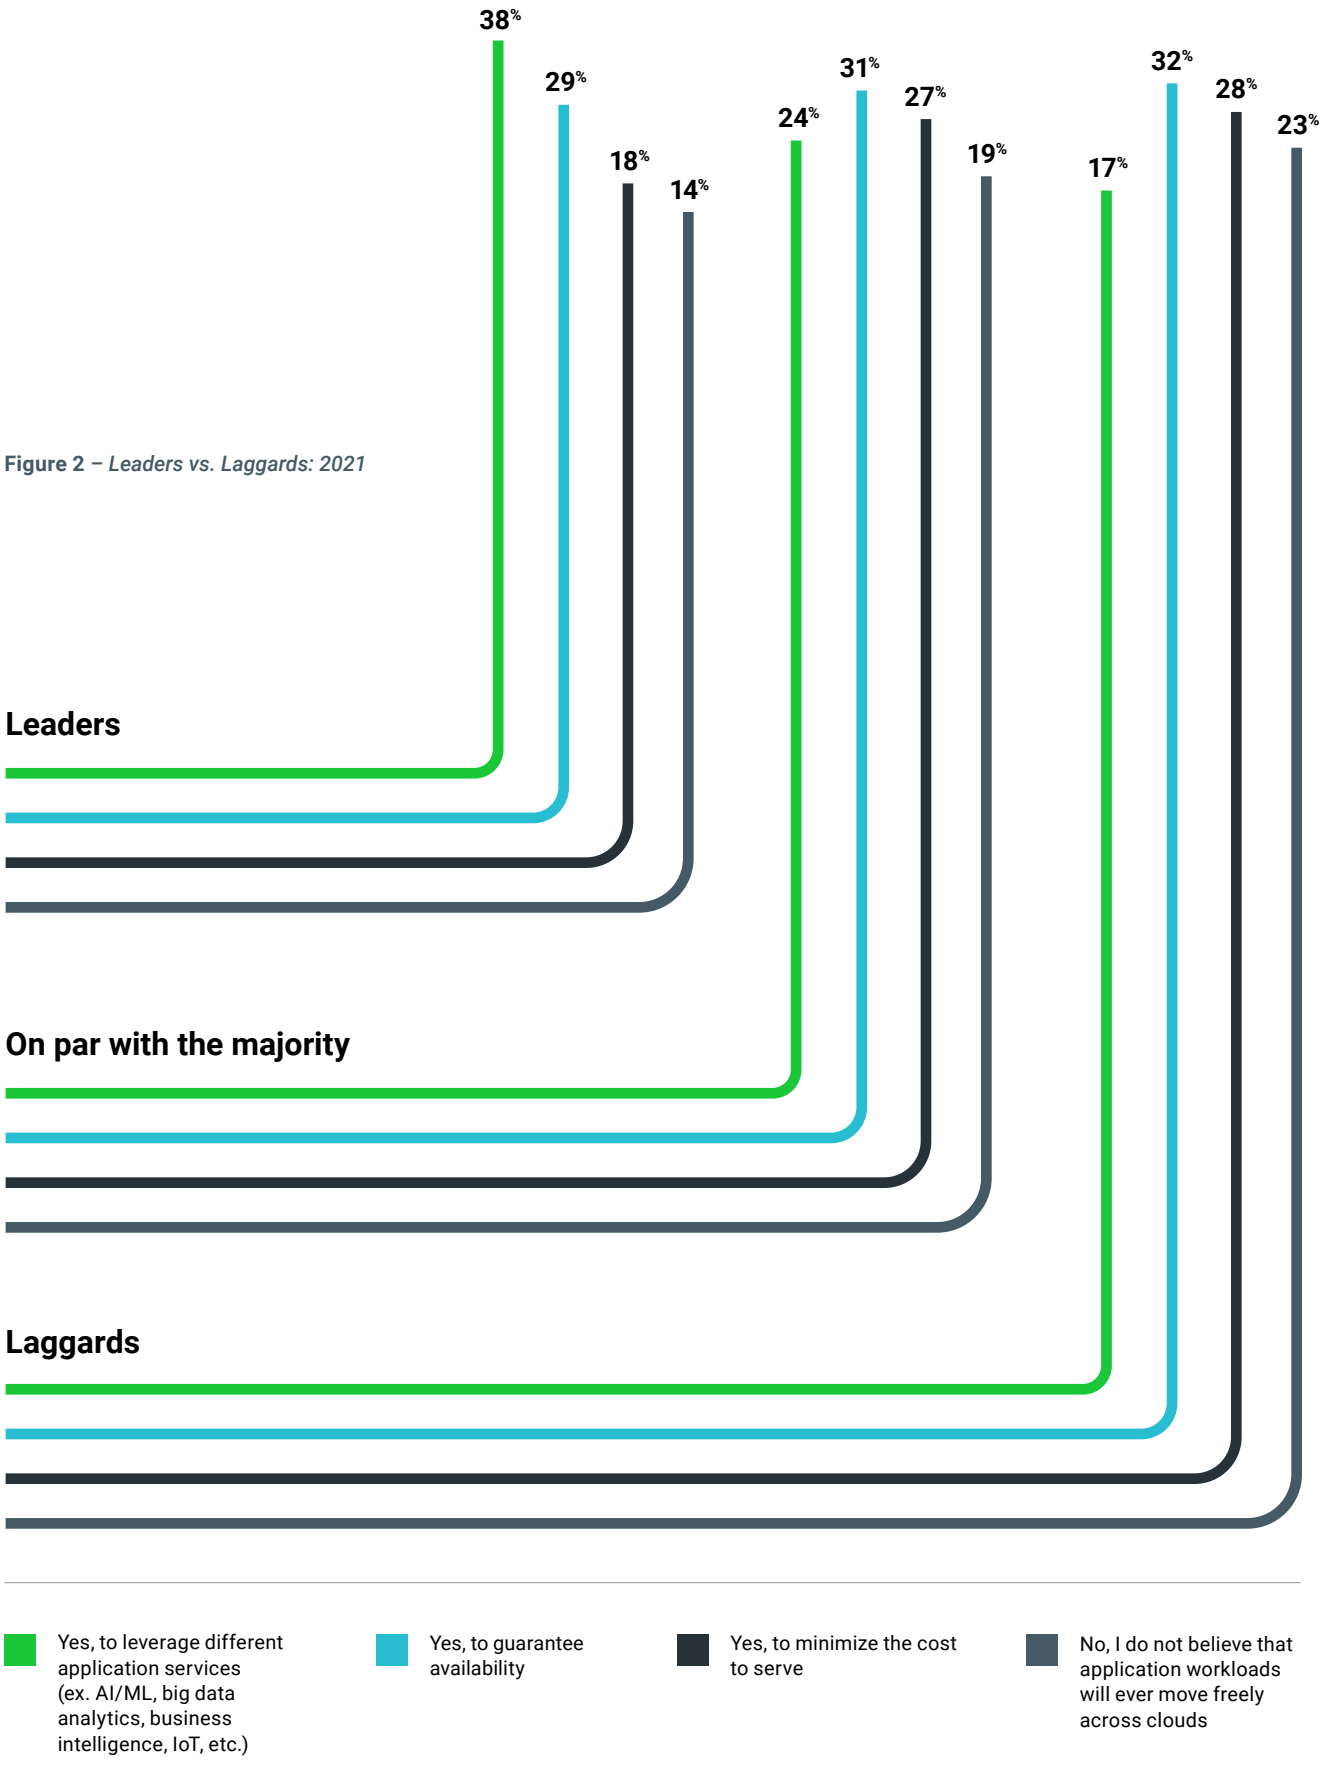

In 2021, leaders value business leverage over the benefits of any one cloud.

In 2021, business leverage takes center stage for leaders. For all the benefits that cloud providers can offer, leading organizations for the first time value "No vendor lock-in" more than "Truly abstracted infrastructure (No patching, updating, etc.)" True, leaders focus on how technology can help them differentiate their business, both in terms of the applications and services they deliver to their customers, as well as the speed to which these competitive advantages can be brought to market. But it is not without its tradeoffs. Architecting applications to leverage specific cloud services naturally creates lock-in. This year the realities of this tradeoff come to bear. As more organizations leverage multiple clouds, we expect internal digital and finance teams will have to navigate these tradeoffs together. After all, clouds will only compete when we can "vote with our feet."

Question: What's more important to you? (2021: N = 819, 2020: N = 938; 2019: N = 846)

Figure 3 – *Overall: 2019-2021*

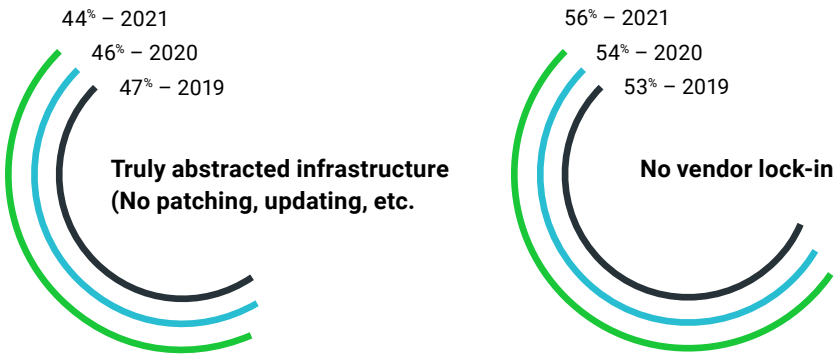

Figure 4 – *Leaders vs. Laggards: 2020-2021 (2021: Leaders - N = 279; Majority - N = 432; Laggards - N = 108*

2020: Leaders - N = 311; Majority - N = 483; Laggards - N = 144)

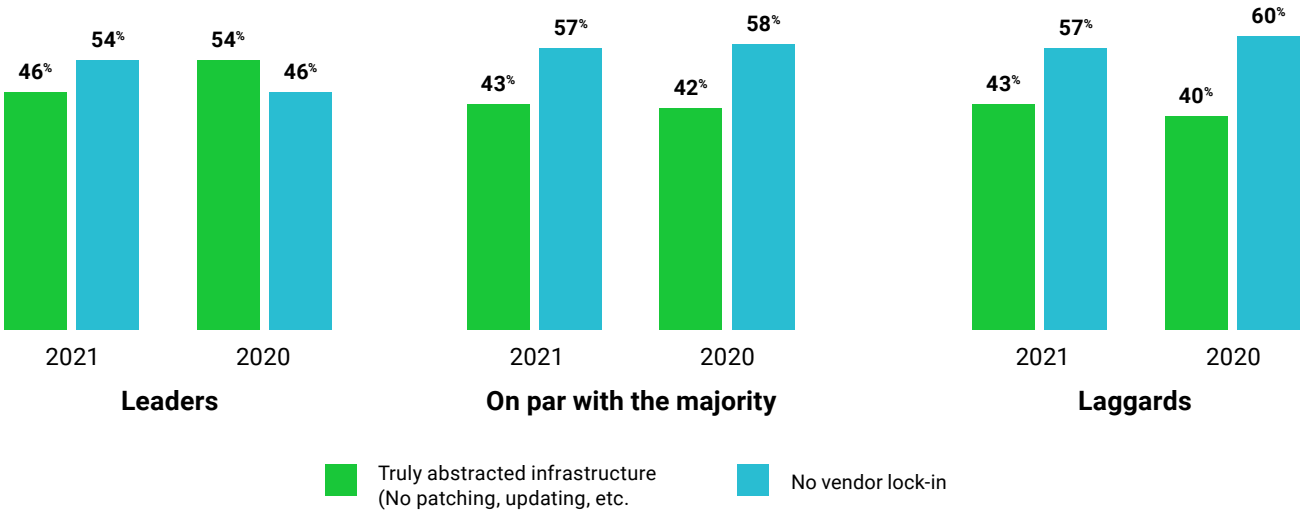

In 2021, leaders prioritize launching new internal services.

For the fourth year in a row, we have asked respondents what they would do if given 30% of their time back. Since 2018 organizations overall have prioritized modifying or improving internal processes and self-education. Launching new services, whether internal or external has always been less of a focus, especially the latter. However, in 2021, leaders have had a notable change of heart. Last year, their priorities were in line with their peers. This year, launching new or improving internal services (ex. self-service app deployment) takes significantly more priority.

To some extent there is an order of operations to these initiatives. An organization cannot launch new services without first ensuring teams have the right skills and operational methods in place. And launching new or improving internal services to, for example, enable developers to work more effectively needs to happen before you can consistently launch new or improve services for customers externally.

Anecdotally, we see organizations in the earlier stages of their transformation who face an uphill battle as they adopt cloud simply because their operational processes, such as ITIL, are outdated. They must reconcile governance with elasticity and agility. Once they do so, they can start solving problems for the business.

Another thing to consider is that the Coronavirus pandemic forced many organizations to adapt quickly. Digital wasn't just augmenting real-world engagement with a brand. In most industries it's become the primary, if not only, form of engagement. Adapting to these dramatic changes in consumer behavior requires new applications and services, for example purchasing online for curb-side or in-store pickup.

Question: If you had 30% of your time back, which of the following would you prioritize?

Figure 5 – *Overall: 2018-2021*
(2021: N = 819; 2020: N = 938; 2019: N = 736; 2018: N = 894)

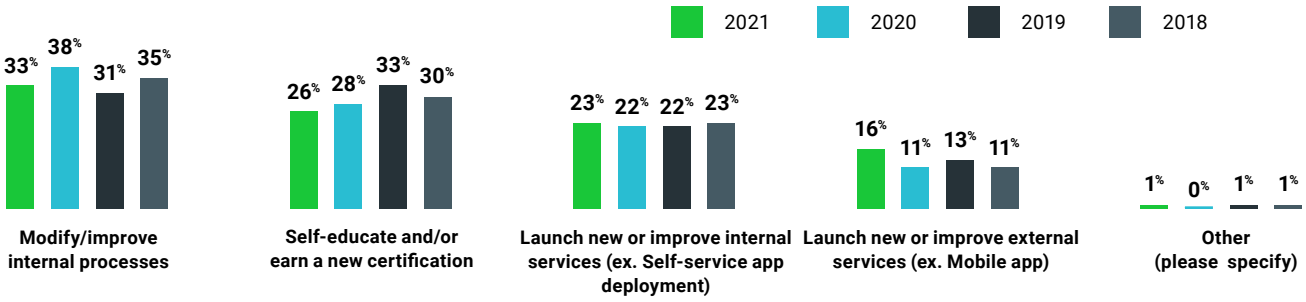

Figure 6 – *Leaders vs. Laggards: 2020-2021*
(2021: Leaders - N = 279; Majority - N = 432; Laggards - N = 108; 2020: Leaders - N = 311; Majority - N = 483; Laggards - N = 144)

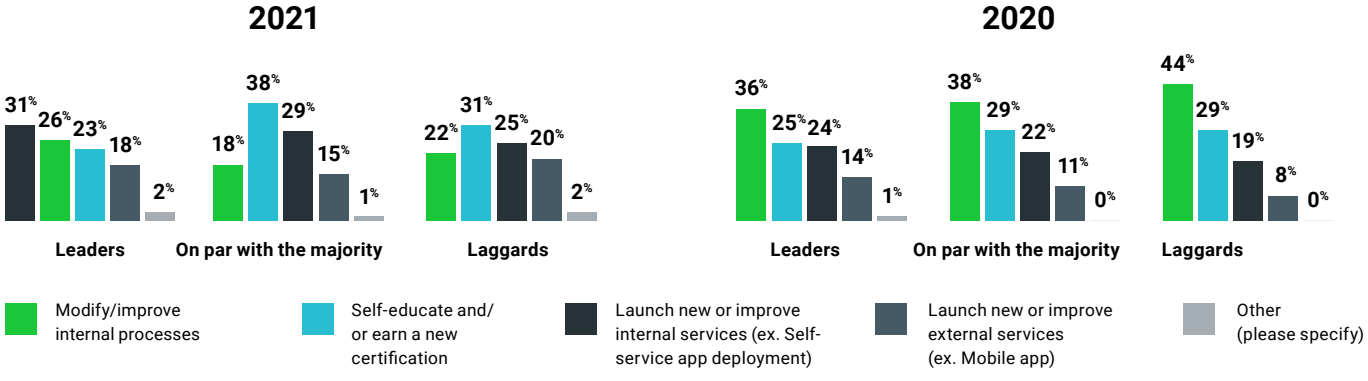

Data gravity continues to be the greatest barrier to dynamic multicloud, followed by security.

We are collecting and using increasingly vast amounts of data in our applications. This year, again, respondents told us that the leading barrier to freely moving workloads is the inability for application components, such as database services, to be fully portable (22%), followed by security (18%).

Data gravity takes two forms: the architectural investment that comes with the platform or services used to store it and make it consumable and the volume of data to be hosted and/or processed in any given location. Both have very real implications for organizations implementing multicloud. The former raises the question of lock-in; the latter comes down to tradeoffs forced by physics (**see Chapter 4: Edge Computing**).

Again, lock-in is a leading concern. While the database-as-a-service offerings themselves may be similar, there are differences in the way cloud providers manage networks, security, and other configurations that do not exactly line up, creating some risk to moving data. Managing data is not a single issue: developers must consider requirements to share data, how to handle replication, and not locking themselves into configurations that could make switching services or leveraging different cloud services difficult. These are tough constraints to navigate, but ultimately people are going to do the best they can to achieve portability, avoid lock-in, and maintain leverage.

Question: What do you believe is the greatest barrier to applications workloads moving freely across deployment platforms (data centers, clouds, edge computing, etc.)? (2021: N = 819; 2020: N = 938)

Figure 7 – *Overall: 2021*

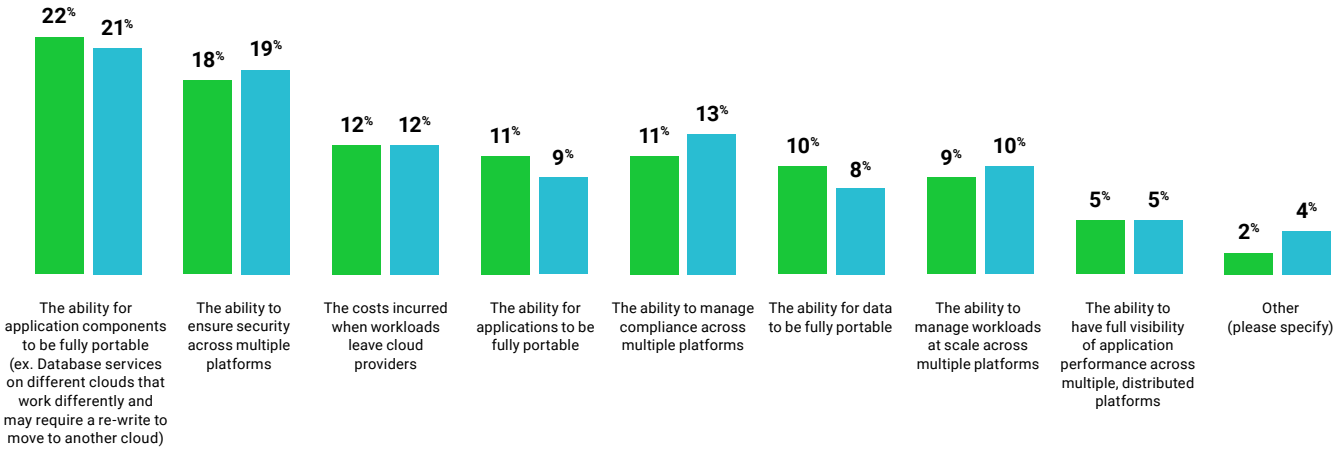

Complexity continues to be the leading challenge for leaders, followed by cultural change.

For organizations as a whole, culture and complexity continue to be the leading (almost tied) challenges for organizations since 2019. While we ask the question generally in our survey, "My organization's top challenges to achieving its goals are...," the results are also consistent with the **November 2020 Cloud Native Survey**, which asks about the challenges faced in using/deploying containers: complexity and culture tied for first place, each with 41%.

Organizations are transforming their businesses through digital applications and services. Cloud and containers are key enablers of this transformation, but they also create complexity and necessitate new modes of operation.

For leaders, specifically, complexity (58%) continues to be their biggest challenge, indicating that cultural change is a prerequisite for successful transformation. Teams must dynamically adapt to rapidly shifting customer behaviors and market trends to stay competitive. As we noted earlier, it requires new modes of operation, namely tighter collaboration between teams, common KPIs and perspectives, agility, and the like.

A promising shift this year is that complexity also became the leading challenge for those on par with the majority (58%) and for laggards it is "catching up," taking second place (55%) after culture. It suggests that organizations are successfully overcoming the other challenges. Did we mention we're all on this journey?

Notably, in 2021, inability to hire or train people with the necessary skill sets is less of an issue (34%)....all that investment in self-education is paying off (**see fig. 5**). We also wonder if the normalization of remote work in the last year has allowed organizations to expand their search for talent.

For the first time, addressing the skills gap has been displaced by the challenge of choosing the right tools and/or platforms (39%), indicating that people are now seriously trying to evaluate solutions. Though real investments in time and money are being made, the cloud and cloud native landscapes are still forming. Early markets will always exhibit fragmentation: everyone gets a shot at trying to solve the problems that must be solved. But ultimately there will be winners and losers—and consolidation.

Question: My organization's top challenges to achieving its goals are: (Pick up to 3) (N = 819)

Figure 8 – *Overall: 2019-2021*

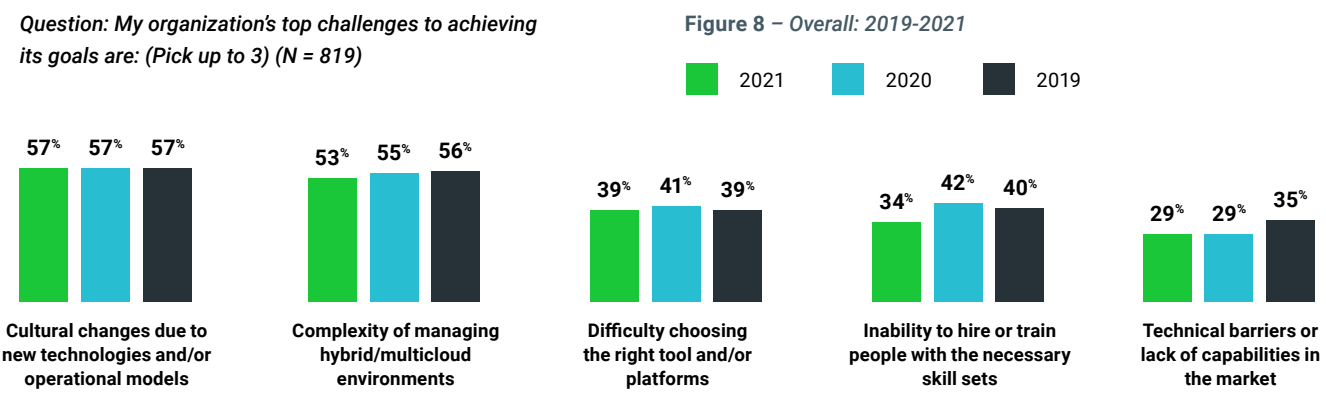

Figure 9 – *Leaders vs Laggards: 2020-2021*

(2021: Leaders - N = 279; Majority - N = 432; Laggards - N = 108; 2020: Leaders - N = 311; Majority - N = 483; Laggards - N = 144)

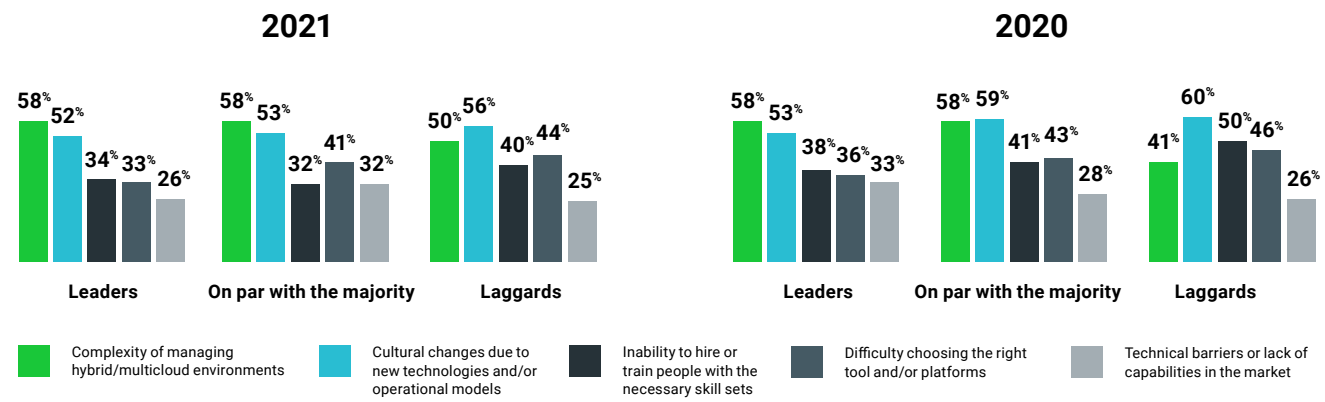

CHAPTER 2:

CONTAINERS

For 61% of organizations containerization will play a strategic role within 18 months; today it is already strategic for nearly 20%.

Containers are the building blocks of the modern application. Lightweight and ready for business in minutes they promise developer speed and elasticity. And their portability is a key enabler of multicloud strategies. So it should come as no surprise that organizations are making the investment. But to see that so many organizations are on track for containerization to play a strategic role tells us that containers have become the standard for digital transformation.

Of course when we look at the data comparing leaders versus laggards, the differences are even more stark. For 71% of leaders, containerization will play a strategic role within 18 months. For the majority and laggards, it is 58% and 48% respectively. Another way to look at it: for nearly half of all self-described laggards containerization will still be strategic within 18 months.

Question: By when do you expect containerization will play a strategic role for your organization?

- Containerization is playing a strategic role today
- Within 6 months
- Within 12 months
- Within 18 months
- More than 18 months from now
- Containerization will never play a strategic role for my organization

Figure 10 – *Overall: 2021 (N = 819)*

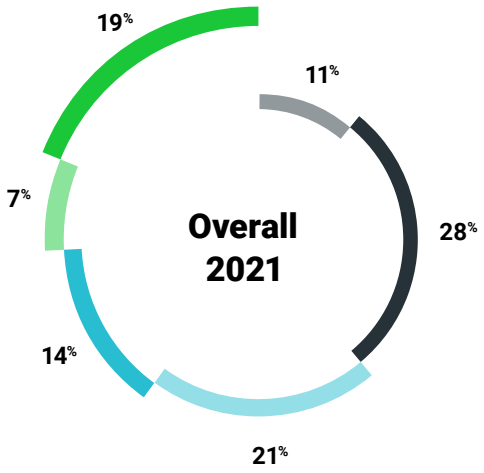

Figure 11 – *Leaders vs. Laggards: 2021 (Leaders - N = 279; Majority - N = 432; Laggards - N = 108)*

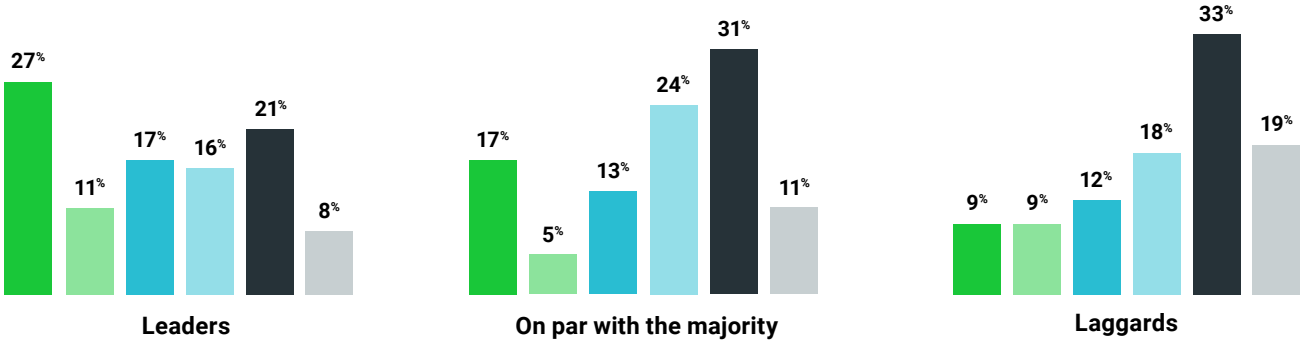

56% are using at least one container platform today, 79% of which are using commercial Kubernetes

We asked respondents to indicate what container platforms they are using today or planning to use within 18 months. Four hundred and fifty nine (459) of the 819 total respondents indicated current usage of at least one container platform (56%). In 2020, 501 of 938 total respondents indicated current usage of at least one container platform (53%). The evidence of more than one container platform suggests that one does not fit all application and infrastructure needs.

Kubernetes claimed victory as the de facto container platform years ago, but 79% of its adoption being commercial indicates that people are putting real money and resources behind it. It is more evidence that this technology is a strategic investment.

Question: Is your organization making a strategic investment in the following container platforms? If not strategically investing in the platform, please indicate that by skipping the row.

Figure 12 – *2021*

(N = 819)

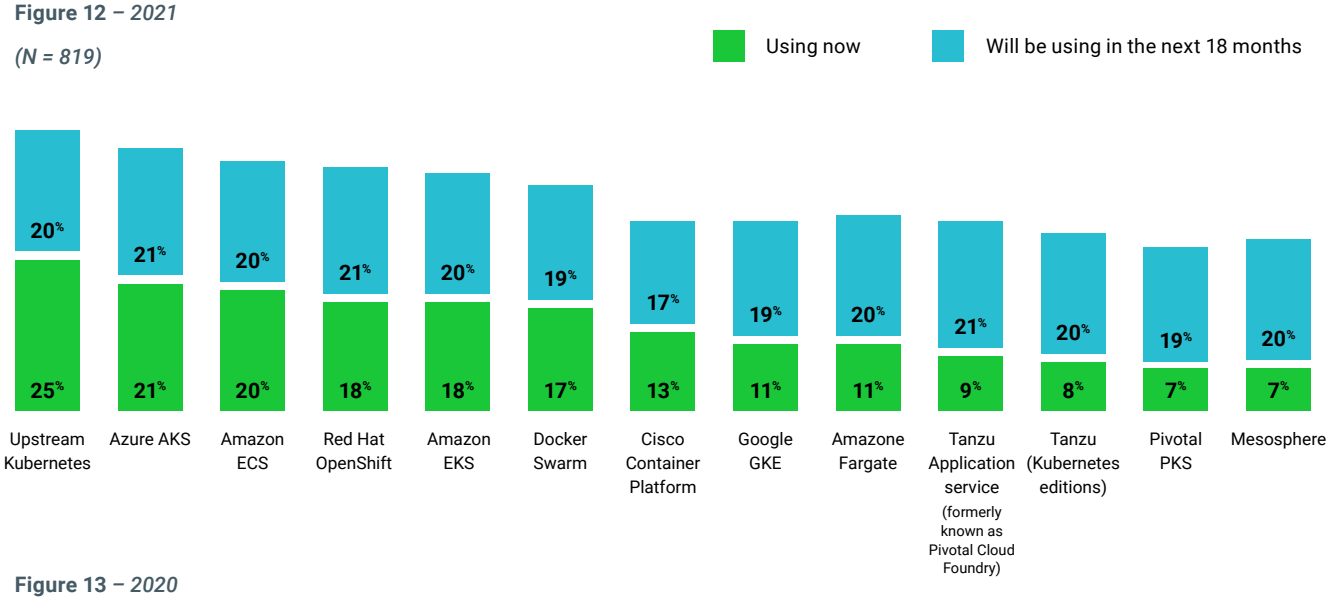

Figure 13 – *2020*

(N = 938)

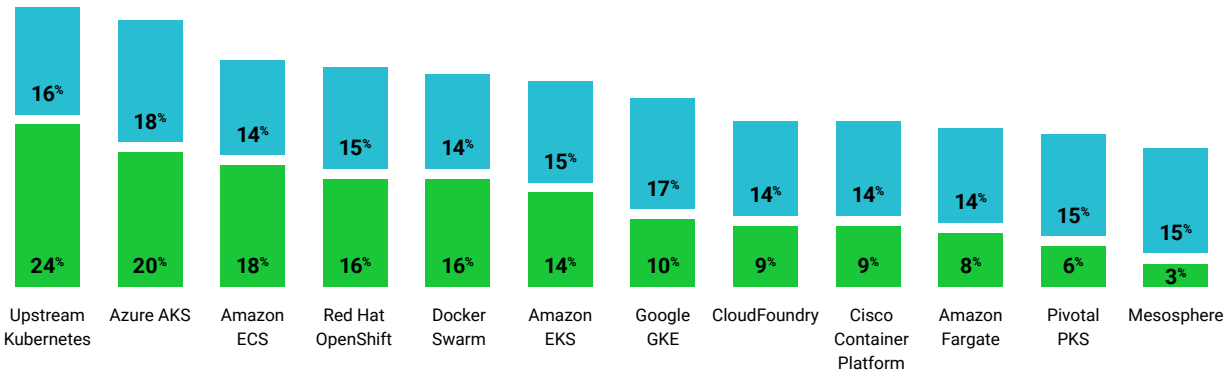

For leaders portability continues to be the leading driver for containerizing applications, but in 2021 efficiency trumps developer speed.

For three years we've asked about the leading drivers behind containerizing applications. For respondents as a whole, portability and efficiency have been the top two drivers, with portability taking the lead this year and last. However, when we look at the responses of leaders versus their peers we see some interesting changes in 2021.

Question: What is the primary reason your organization is containerizing applications? Please skip if you're not using containers. (2021: N = 664/819; 2020: N = 685/938; 2019: N=494/846)

Figure 14 – *Overall: 2019-2021*
(2021: N = 664/819, 2020: N = 685/938; 2019: N = 494/846)

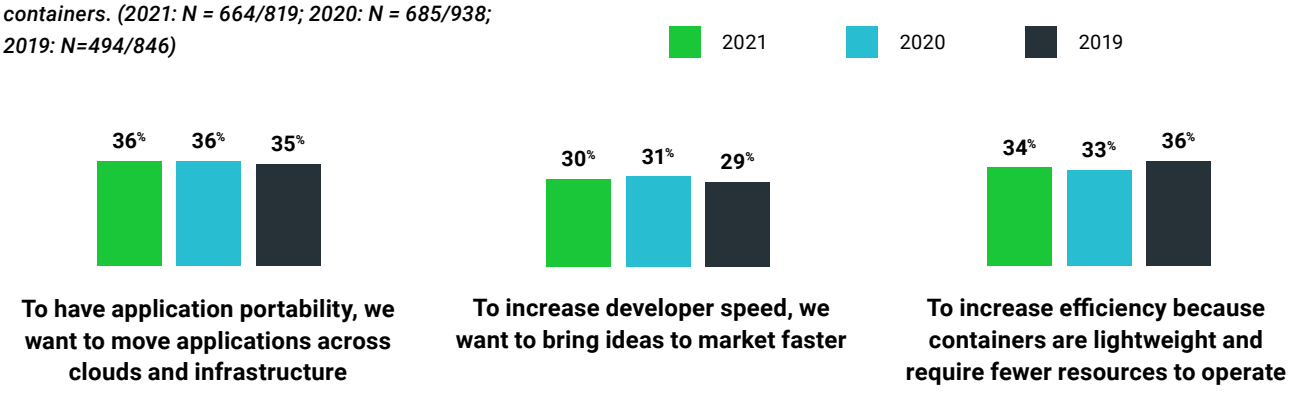

In 2020 leaders were motivated by portability (37%), then to increase developer speed (35%), and, lastly, efficiency (28%). This year portability was still the primary driver (38%), but efficiency (32%) took a slight lead over developer speed (30%). For those on par with the majority, efficiency took the lead (38%) this year; last year it came second (35%) to portability (37%). Meanwhile, laggards are motivated by developer speed (39%), portability (32%), and, lastly, efficiency (29%).

We reckon that the fact that leaders and those on par with the majority are giving more consideration to efficiency is a symptom of both the fact that adoption has increased and it is costly to maintain an overprovisioned environment, along with a rapid adoption of public cloud which provides infrastructure elasticity, but at a price. After all, to really reap the benefits of pay-as-you-go cloud resources you have to architect your applications to be able to scale out or back based on demand. That's true elasticity, enabled by containerization. Without containers, applications are more often sized to peak, as they cannot scale to peak.

Figure 15 – *Leaders vs Laggards: 2020-2021*
(2021: Leaders - N = 248; Majority - N = 340; Laggards - N = 76
2020: Leaders - N = 257; Majority - N = 349; Laggards - N = 79)

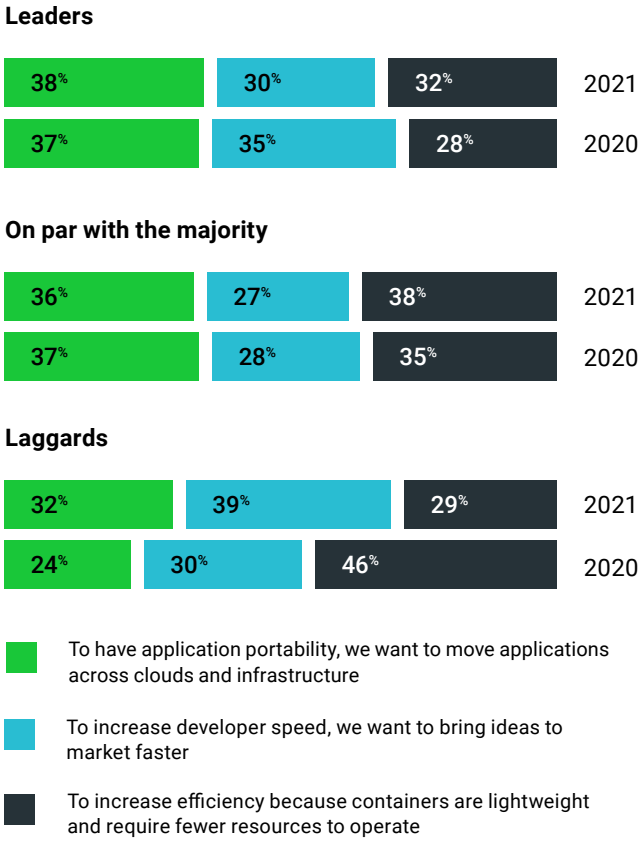

56% of those on their container/cloud native journey are running containers in production, but complexity at scale is hindering advanced use cases.

Since 2019 we've asked respondents to share where their organization is on its journey to containers/cloud native. We are seeing incremental maturation on this journey. In 2021 slightly more organizations on their containers/cloud native journey are in production, 56% up from 53% last year and 52% the year before. We see the most growth in Early Production, 33% up from 27% last year. However, those in Advanced Production or Platform-First remain flat, suggesting perhaps that the operational realities of managing containerized applications at scale stalls further advancement for most organizations.

Question: Where is your organization in its journey to containers/cloud native?

Figure 16 – *Overall: 2019-2021*

(2021: N = 598/819, 2020: N = 701/938; 2019: N = 527/846)

Note: Percentages shown here are of those that are on the journey (ex. For 2021, 598/819 respondents are on their journey and 28% of them are exploring). Respondents that answered "I don't know" or "N/A, we are not using containers" are not shown in this graph.

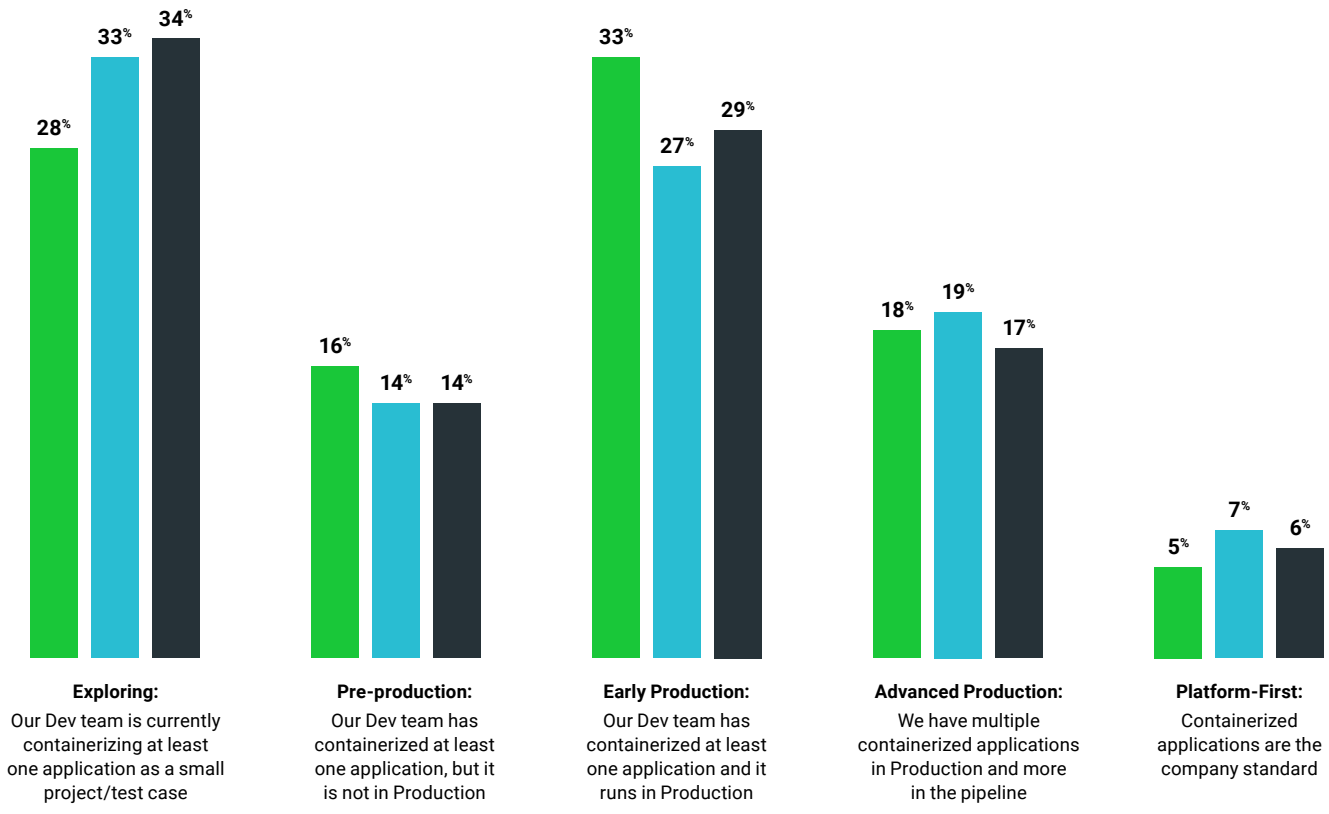

As previously discussed, complexity remains the biggest challenge for leaders. And, specifically with regards to containers, the **2020 CNCF Survey** found that complexity is the leading challenge to using/deploying containers.

Looking at net new adoption trends, in 2021, 73% of organizations are on this journey, up from 62% in 2019, but flat since 2020.[1] When we asked how the pandemic affected their organization's digital priorities, only 9% of respondents initiated building new microservice/cloud native applications and only 7% initiated refactoring existing applications.

Question: How has the Coronavirus pandemic affected your organization's current digital priorities?

Figure 17 – *Overall: 2019-2021*

(N = 819)

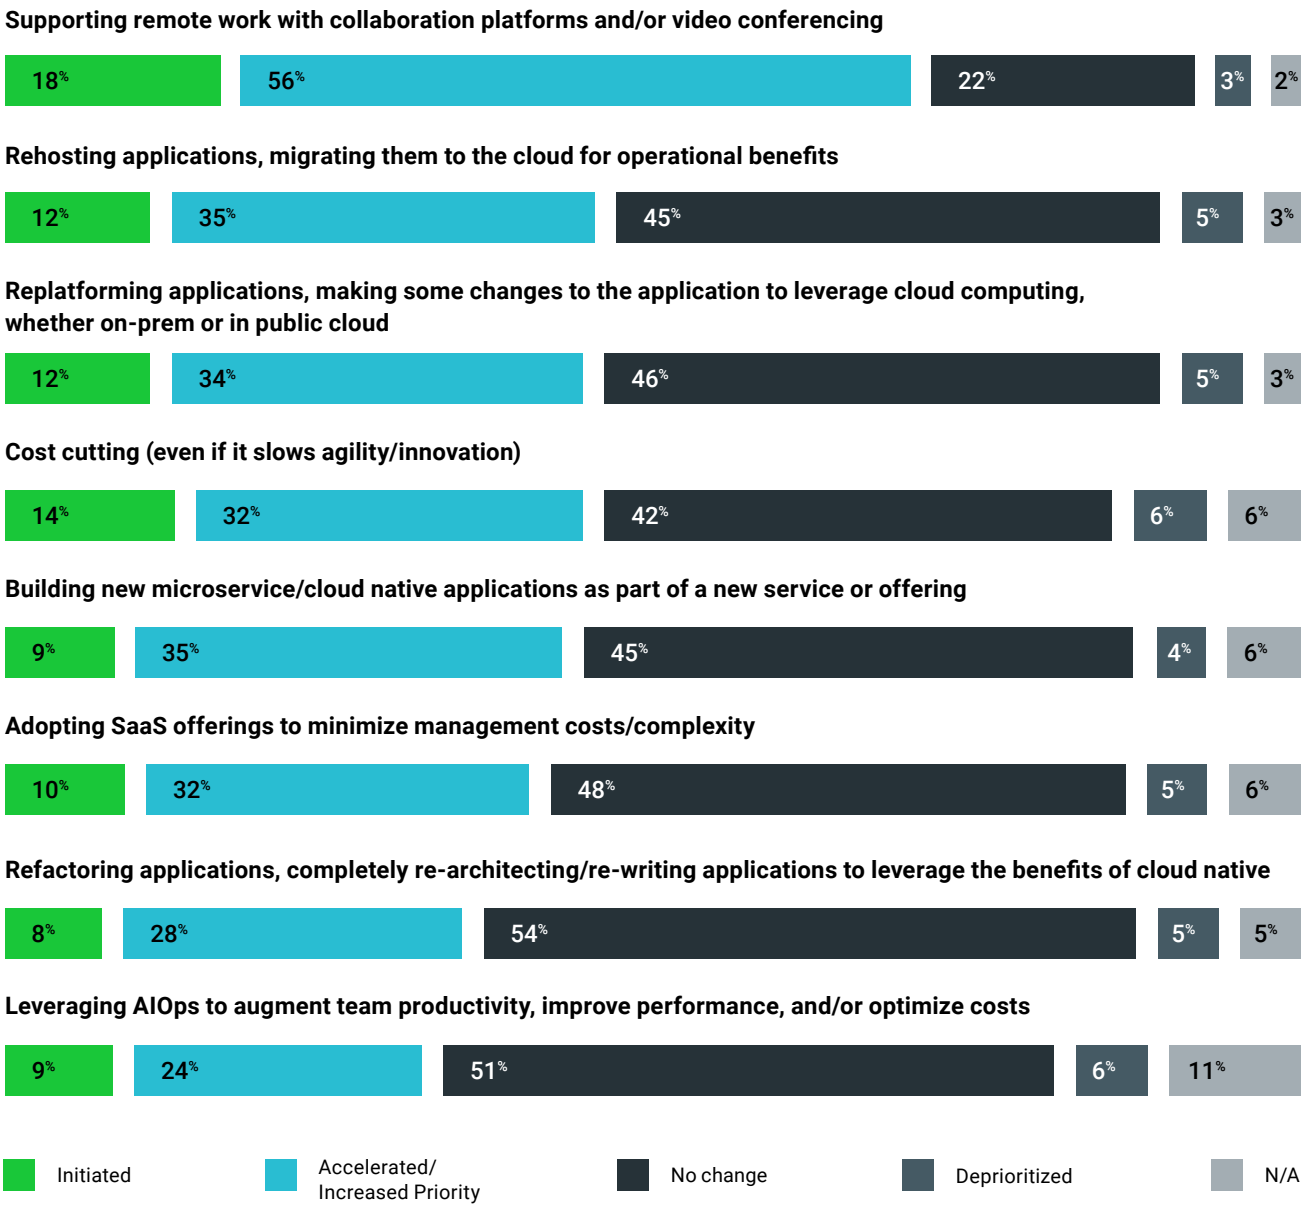

[1] 598 out of 819 respondents placed their organization somewhere on this journey i.e. did not choose "N/A" or "I don't know." 598/819 = 73%. Likewise, these percentages were calculated for 2020 and 2019: 701/938 = 75% and 527/846 = 62%, respectively.

54% are using at least one container platform add-on, up from 48% in 2020.

We asked respondents to indicate what container platform add-ons they are using today or planning to use within 18 months. Four hundred forty-two (442) of the 819 total respondents indicated current usage of at least one container platform add-on (54%). In 2020, 448 of 938 total respondents indicated current usage of at least one container platform add-on (48%).

The use of messaging with container platforms has increased to 21%, bumping up to third place. It validates that organizations are indeed building applications with multiple services, which need to communicate with other services or applications. Hello, microservices.

Question: Is your organization making a strategic investment in the following container platform add-ons? If not strategically investing in the add-on, please indicate that by skipping the row.

Figure 18

2021: N=550/819

2020: N= 549/938

Note: Percentages shown here are of those that selected at least one option.

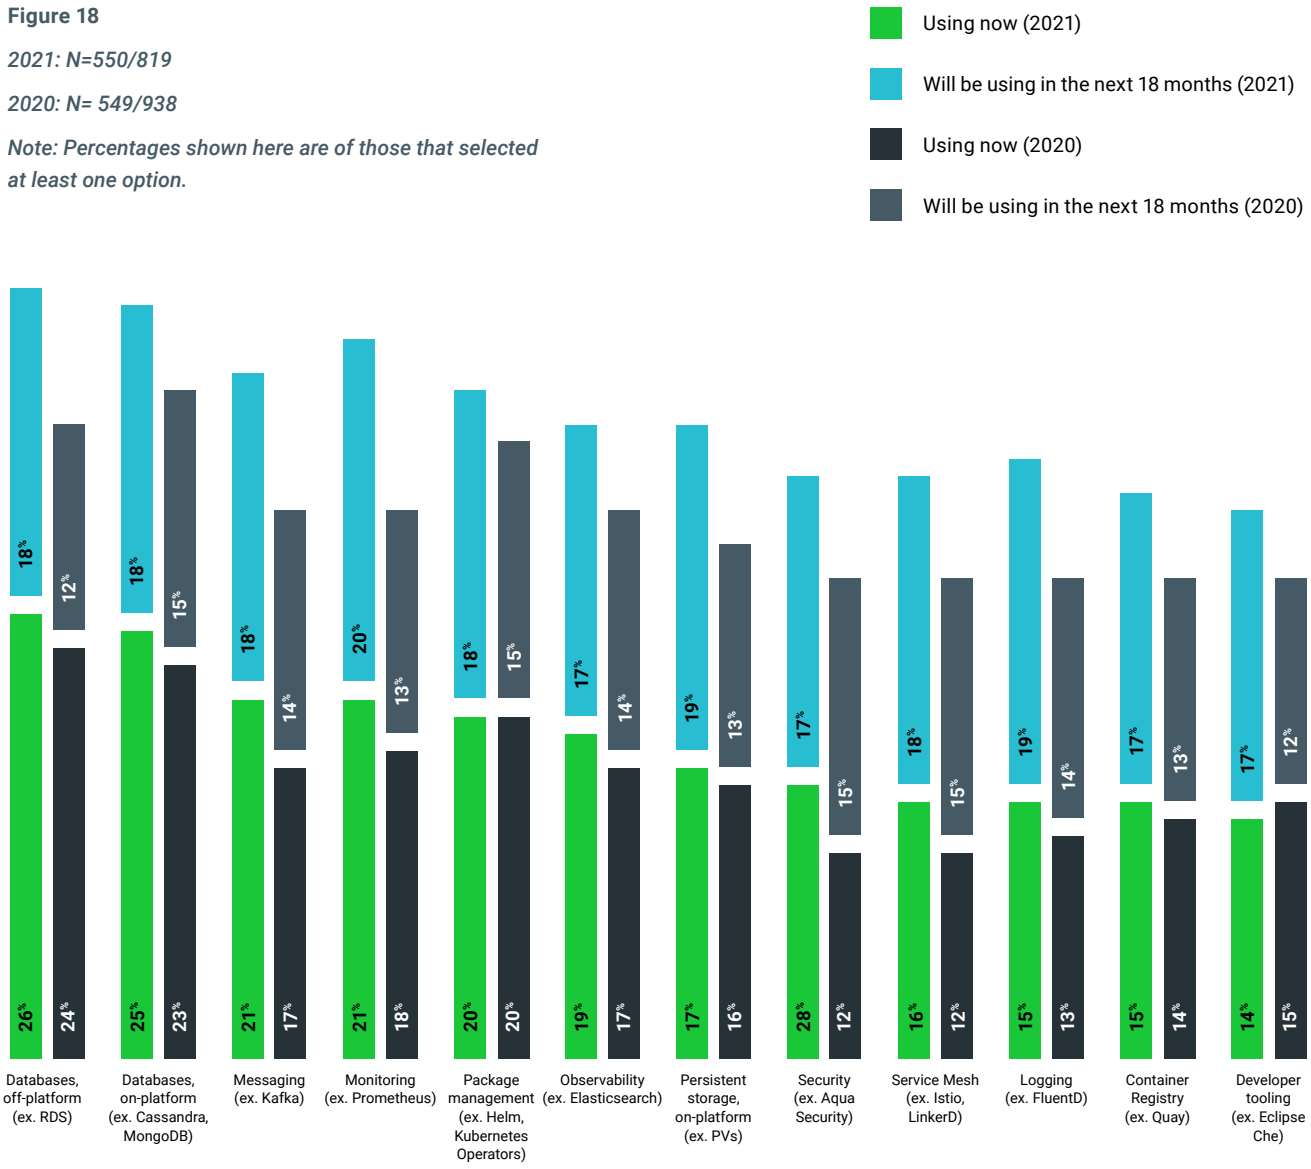

In 2021 containers running in public cloud took the lead over on-prem and the trend suggests continued growth in public cloud.

One of the key benefits of Kubernetes—the leading container platform—is that it can run anywhere. Looking at respondents overall, this was the first year that containers running on public cloud infrastructure (56%) took a lead over containers running in a virtualized on-premises environment (53%), in 2020 those numbers were 51% and 56% respectively.

While apps and services differentiate cloud providers, public cloud also affords its customers more elasticity, if they have the applications that can take advantage of it. As we noted earlier, one of the key benefits of containerization is that it allows individual services of an application to scale out and back on demand. True elasticity. Public cloud infrastructure can support this elasticity, spinning up Kubernetes nodes as needed, spinning them down when not needed; only paying for the infrastructure you use.

On-premises infrastructure is limited in that spinning up nodes requires the assurance that the capacity, right down to the hardware, is available. An organization has to maintain some buffer of capacity in anticipation of rapid growth, as well as scaling to peak.

The elasticity of containerized applications is bringing the business benefits of public cloud OpEx versus traditional CapEx models to bear.

Question: On what type of infrastructure are you running your container platform(s)? Select all that apply. Please skip, if not using containers.

Note: Percentages shown are of those that answered the question. Respondents were asked to skip this question if they are not using containers.

Figure 19 – *Overall: 2019-2021*

(N=633/819; 2020: N= 667/938; 2019: N = 466/846)

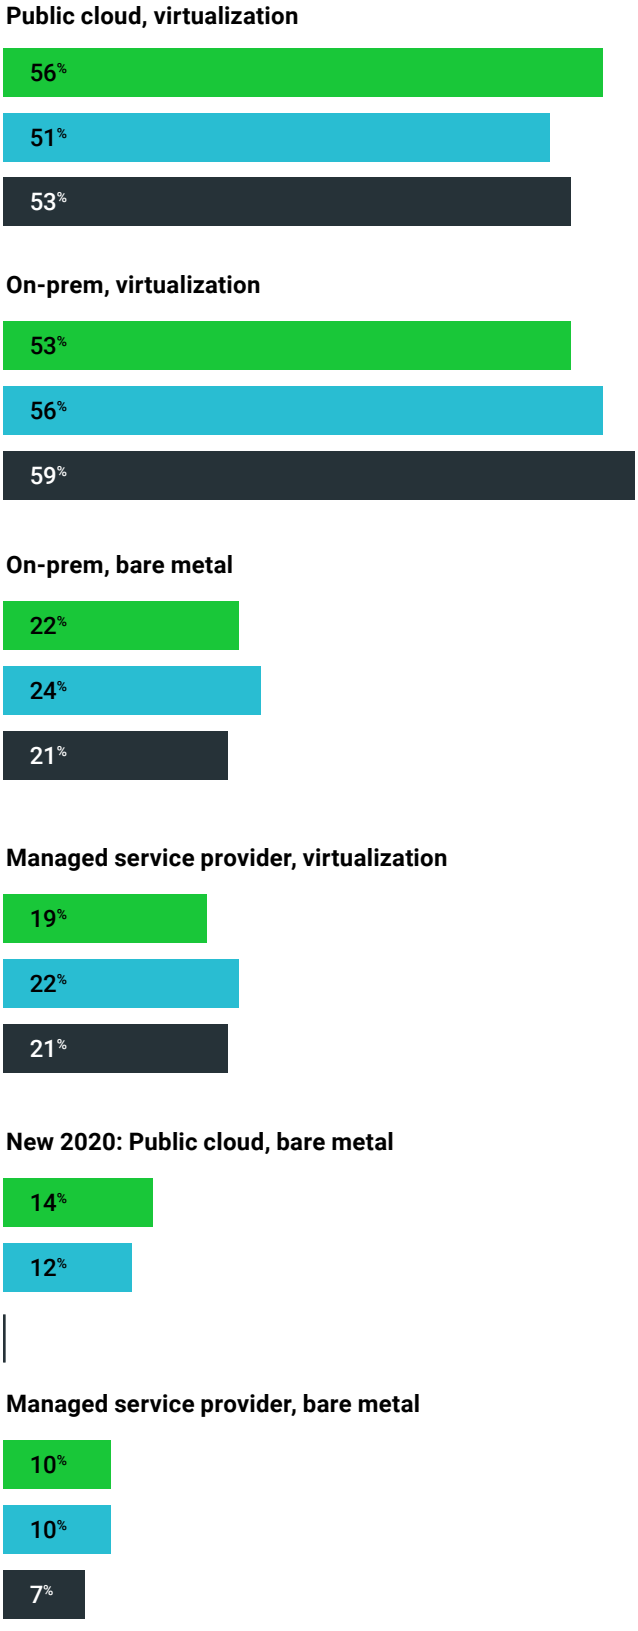

For leaders, containers are overwhelmingly being run in the public cloud (65%) versus on-premises (50%). No doubt this growth off-premises is due to the elasticity that public cloud infrastructure affords, if well-managed, as well as the services that cloud providers offer on top of IaaS.

Figure 20 – *Leaders vs Laggards: 2021*

(Leaders - N = 239; Majority - N = 321; Laggards - N = 73)

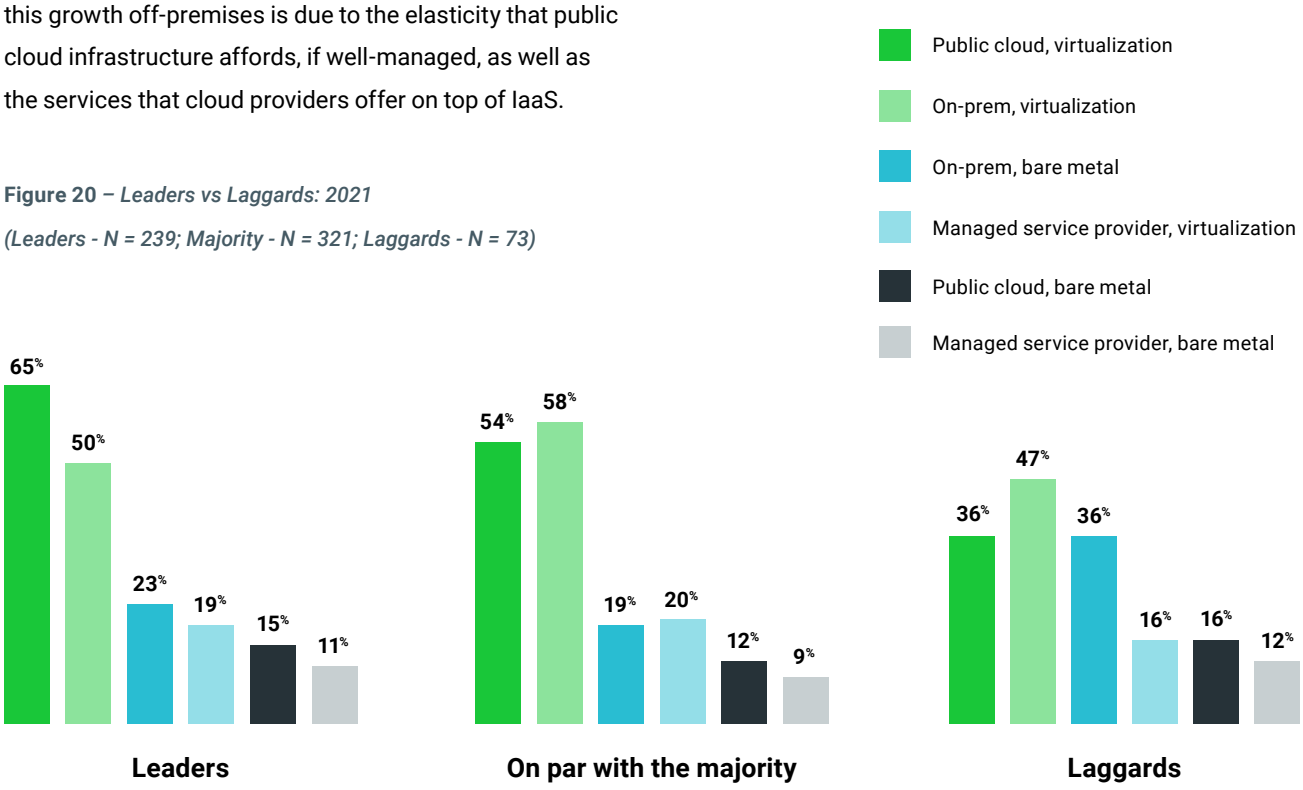

Among those running containers, it's an even split between those running on a single type of infrastructure (50%) versus two or more types of infrastructure. Different infrastructure affords different benefits. While public cloud offers OpEx models that elastic infrastructure can take full advantage of, many organizations will continue to maintain a presence on-premises, be it for security reasons, compliance, or the data gravity of legacy infrastructure. Of course, the more heterogeneous infrastructure, the more complexity.

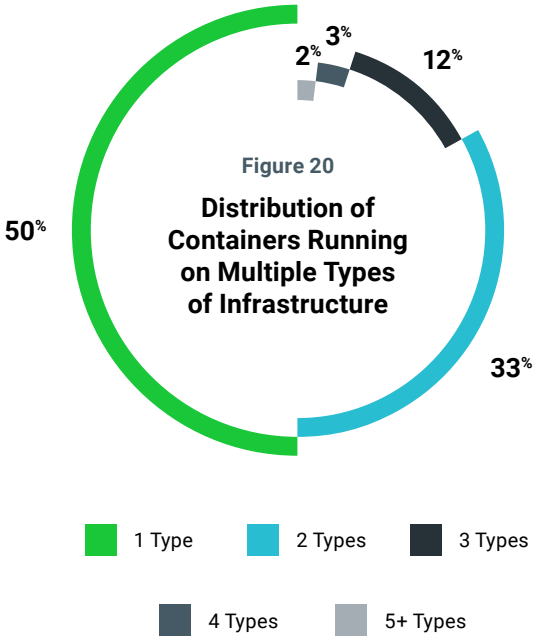

Figure 20

Distribution of Containers Running on Multiple Types of Infrastructure

21% of those running containers are doing so *only* in a public cloud.

Among the 56% running containers in at least one public cloud, 21% are doing so only in public cloud, 21% in a public cloud with one additional type of infrastructure, and 14% in a public cloud with two or more additional types of infrastructure.

Figure 21 – *Breakout of how organizations are leveraging public cloud with other types of infrastructure to run containers.*

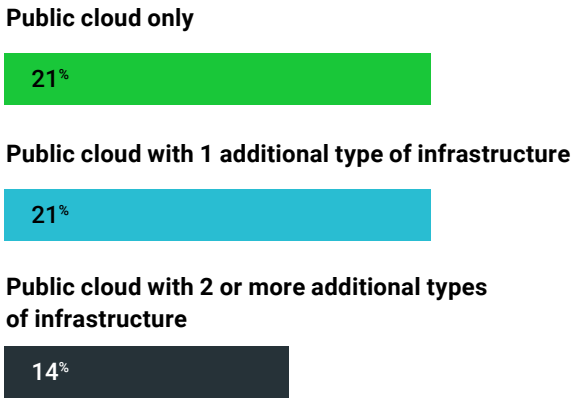

37% of those running containers are doing so on bare metal

Of those that answered the question, "On what type of infrastructure are you running your container platform(s)? Select all that apply. Please skip, if not using containers," 37% are running containers on bare metal.[2] (That's 29% for respondents overall.)

Managing bare metal infrastructure is a different set of skills. But having more experience with Kubernetes at least, which runs anywhere, gives organizations the confidence to explore and possibly take advantage of bare metal's benefits.

What's interesting is how the laggards running containers responded. They, not surprisingly, make up the smallest proportion of those running containers (N = 78).[3]

However, a greater proportion of them are running containers on bare metal: on-prem, bare metal (36%); public cloud, bare metal (16%), and managed service provider, bare metal (12%). To these self-described laggards, can we say, don't be so hard on yourself?

In any case, as we've noted, there are many reasons that compel organizations to maintain workloads on-premises. Once that decision has been made, they may then be motivated to at least cut licensing and operational costs by going bare metal (see next section).

[2] 237 respondents selected at least one of the three bare metal options: on-prem, bare metal; public cloud bare metal; managed service provider, bare metal. 237/633 (the number of respondents that answered this question) = 37%. 237/819 (the number of respondents in total) = 29%

[3] Compared to leaders: N = 239 and majority: N = 321

Performance (42%) continues to be seen as the primary benefit of running containers on bare metal, followed by reducing licensing costs (26%)

Among all respondents, performance is seen as the leading benefit whether for CPU-intensive workloads generally (29%) or access to special hardware (an additional 13%). Reducing licensing costs is the next most popular benefit. Interestingly, there was general agreement between leaders and those on par with the majority, with the exception that the latter was more likely to believe there were no benefits to bare metal over virtualization (18%). For laggards, seeing no benefits came in second (25%) only to reducing licensing costs (28%). Nevertheless, even for laggards, performance combined (34%) is seen as the leading benefit.

Question: What do you see as the primary benefit of bare metal infrastructure for container platforms (over virtualization)?

Figure 22 – *Overall: 2019-2021*
(2021: N = 819, 2020: N = 938; 2019: N = 629)

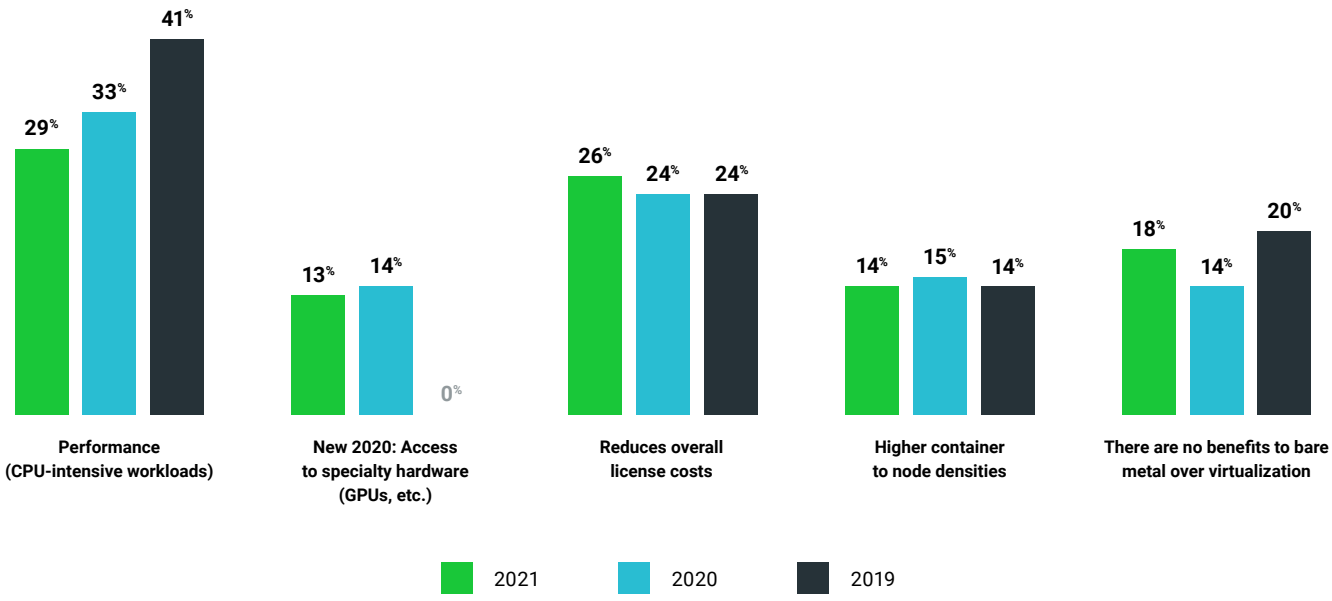

We expect that as organizations continue to build on their experience with Kubernetes, we will see more of them running on bare metal for the performance and reduced licensing benefits. Anecdotally, our own customers are looking to Turbonomic to help them here by way of optimizing their Red Hat OpenShift clusters on bare metal. With bare metal organizations must handle longer cycles to spin up another node, so optimizing the node is even more important. By running OpenShift on bare metal, these customers are looking to avoid the licensing costs of virtualization, as well as the operational/skill-set costs that naturally come with supporting more layers in a stack.

For all these reasons—performance and reduced licensing benefits, as well as growing experience in the container platform that allow them to take advantage of these benefits—half of leaders believe that container platforms will mitigate the need for virtualization (**see fig. 25**).

Figure 23 – *Leaders vs Laggards: 2021*

(2021: Leaders - N = 279; Majority - N = 432; Laggards - N = 108)

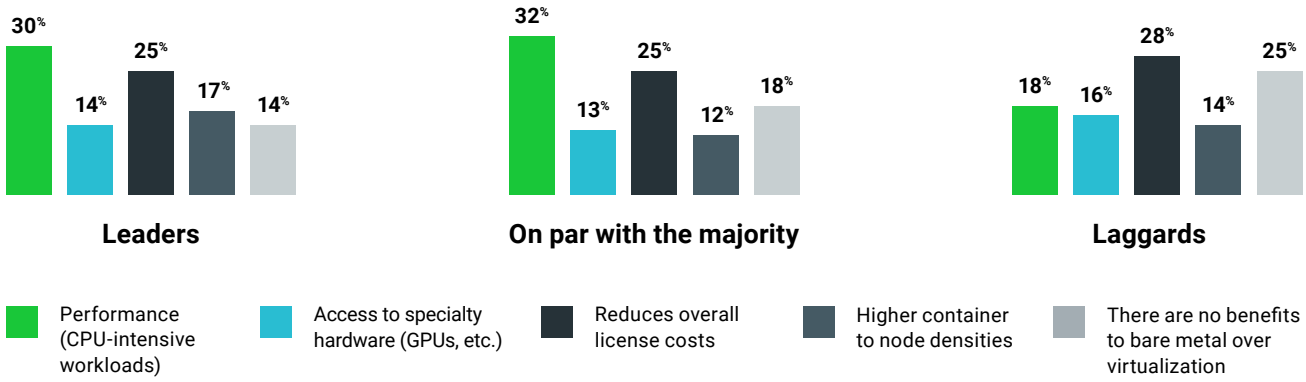

Question: Do you believe that container platforms will someday mitigate the need for virtualization?

Figure 24 – *Overall: 2019-2021*

(2021: N = 819, 2020: N = 938; 2019: N = 684)

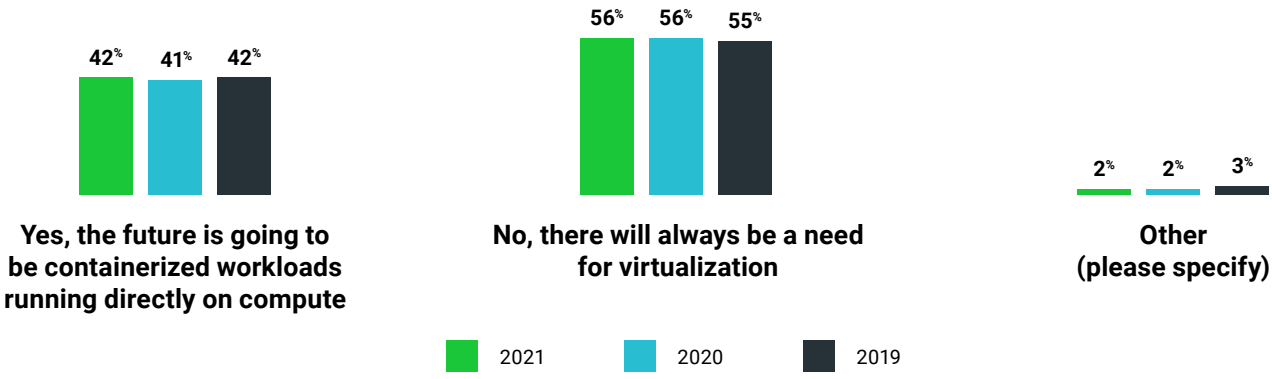

Figure 25 – *Leaders vs Laggards: 2021*

(2021: Leaders - N = 279; Majority - N = 432; Laggards - N = 108)

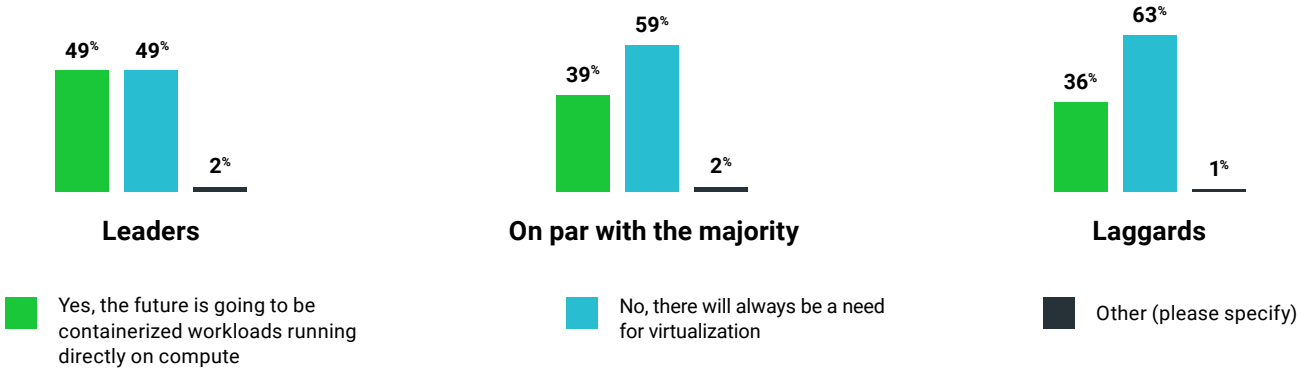

54% of leaders are running stateless applications today, compared to 44% overall.

We asked respondents about their mix of stateless, stateful, and composite applications and found that while 44% of organizations are running stateless applications today, that jumps to 54% for leaders. Respondents also anticipate the most growth in composite applications: 40% expect to have composite applications within 18 months.

As we've previously discussed, containerized applications have the ability to be more elastic. More specifically, it's containerized applications architected with stateless services that deliver this benefit. With this elasticity in the application you are in a position to take advantage of elastic infrastructure. You can build out your Kubernetes clusters for the maximum number of pods that can burst, or you can look at your clusters as something that can be elastic, spinning nodes up and down based on demand. Coupled with the elasticity of public cloud infrastructure, stateless applications allow organizations to realize the full benefits of OpEx models. It explains the growth that respondents anticipate: an additional 32% expect to have stateless applications within 18 months.

Question: What is your mix of stateful vs. stateless (i.e. cloud native) applications? Note: Respondents had to pick at least one of the six options.

Figure 26 – *Overall: 2021*

(N = 819)

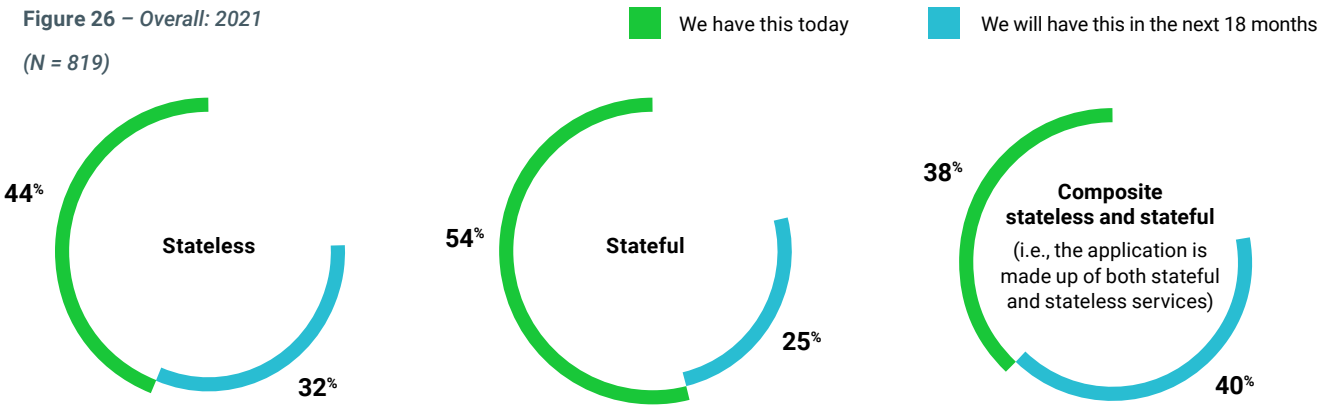

Figure 27 – *Stateless*

(Leaders - N = 279; Majority - N = 432; Laggards - N = 108)

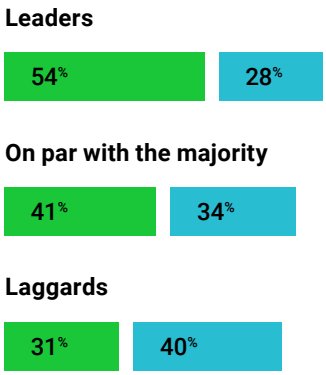

Figure 28 – *Stateful*

(Leaders - N = 279; Majority - N = 432; Laggards - N = 108)

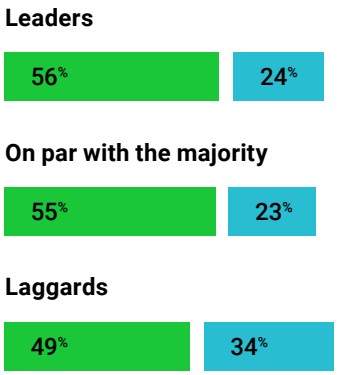

Figure 29 – *Composite stateless and stateful (i.e., the application is made up of both stateful and stateless services)*

(Leaders - N = 279; Majority - N = 432; Laggards - N = 108)

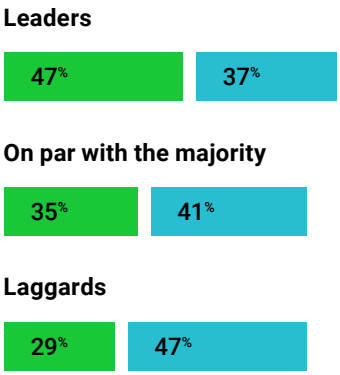

Composite applications will see the most growth, 41% plan to shift or refactor existing stateful applications

While there are clearly benefits of decoupling your application into agile, stateless services, there remains a need to work with persistent data, which could be coming from refactoring classic applications where not everything can be fully stateless.

Respondents expect significant growth in composite applications (those made up of both stateless and stateful services) with 40% saying they will have it in the next 18 months. It suggests that containerized apps and the platforms they run on increasingly need to support both. Organizations see the benefit of containerization for stateful services, further evident by 41% of respondents saying they will shift or refactor, this is unchanged from 2020.

Question: What is your organization's plan for your stateful applications?

Figure 30

(2021: N = 819; 2020: N = 938)

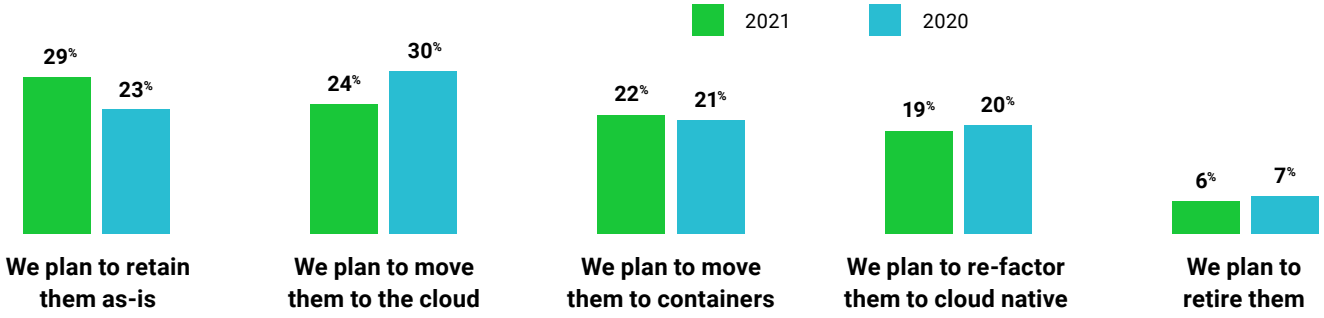

37% believe that handling the resiliency of application data is the biggest challenge to running stateful applications in containers.

Here respondents were asked if their data challenges were more platform and performance based, or concerned with data resiliency. The leading challenge at 35% is resiliency of the application data, suggesting that most do not think that containerized databases and platform services performance is a problem. Application developers need to make choices on how to handle data that does not deter from the agility gains going to microservices.

Figure 31 – *Overall 2021*

(N = 819)

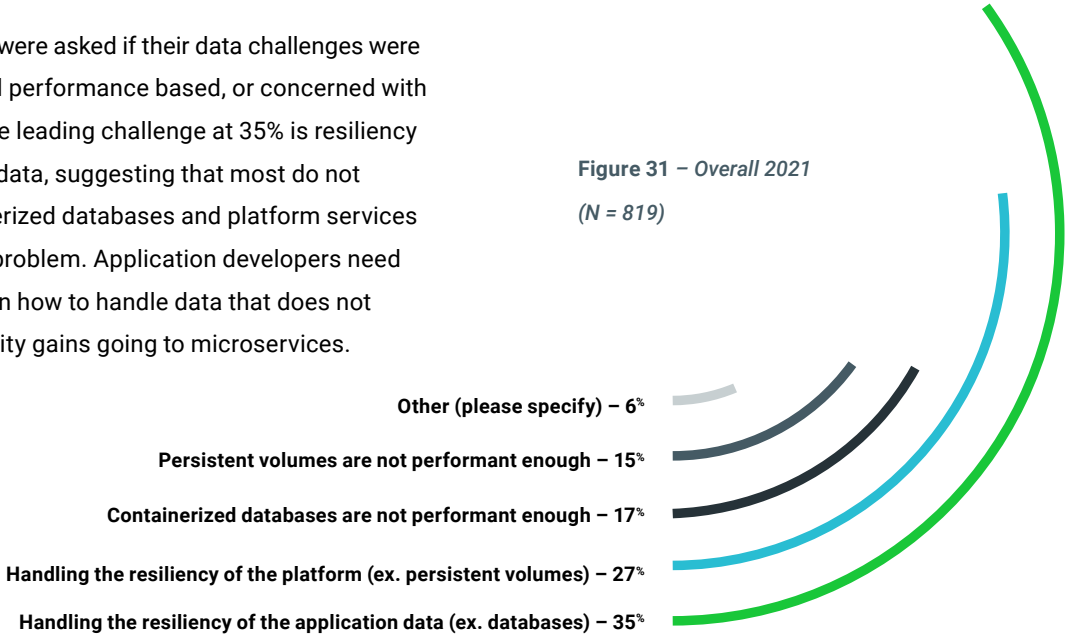

CHAPTER 3:

CLOUDS

Optimization (25%) is the most important initiative for organizations adopting public cloud in the coming year, followed by advancing a multicloud strategy (21%).

With over a decade of experience collectively behind us, organizations are increasing their use of public cloud resources with eyes wide open to the challenges of managing performance versus cost. It is why, in the year ahead, public cloud optimization is the leading most important initiative for 25% of organizations, followed by advancing a multicloud strategy (21%).

For leaders, these two initiatives exhibit a more stark lead. Notably only 12% considered migrating workloads to the public cloud as their most important initiative, indicating that they have already executed this part of their cloud journey. Conversely, laggards and those on par with the majority put cloud migration at 24% (first) and 19% (tied for second) respectively.

Question: Which is the most important initiative your organization will be tackling in the coming year, as it pertains to public cloud adoption? (N = 819)

Figure 32 – *Overall: 2010-2021*

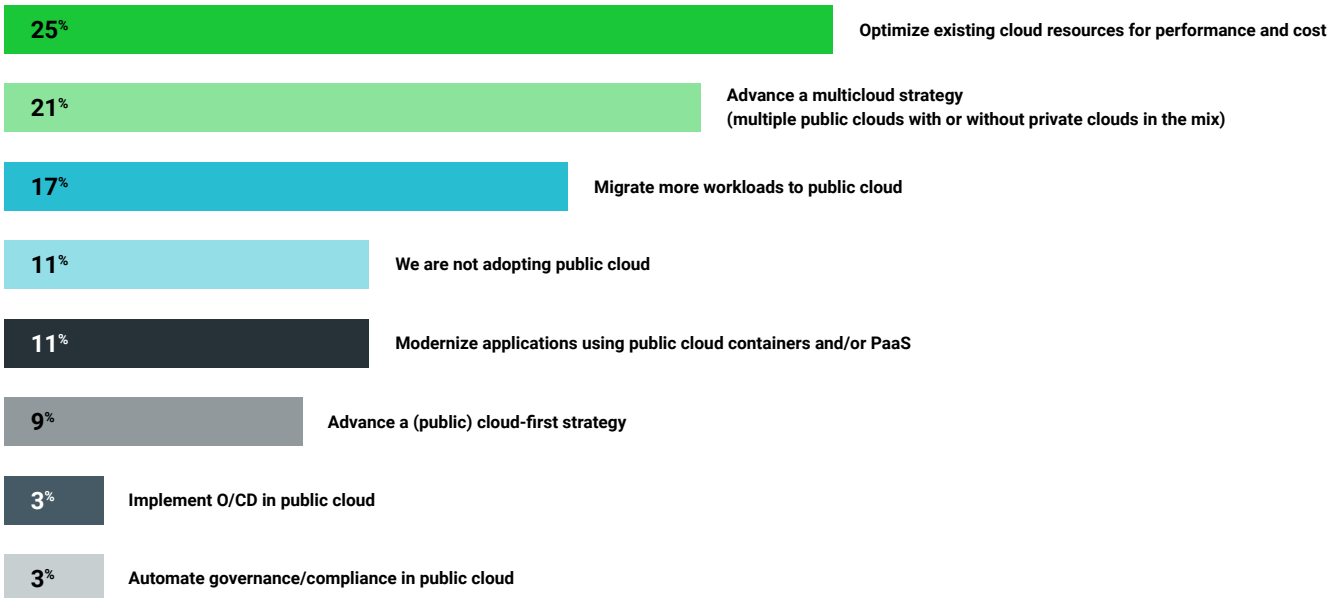

Figure 33 – *Leaders vs Laggards: 2021*

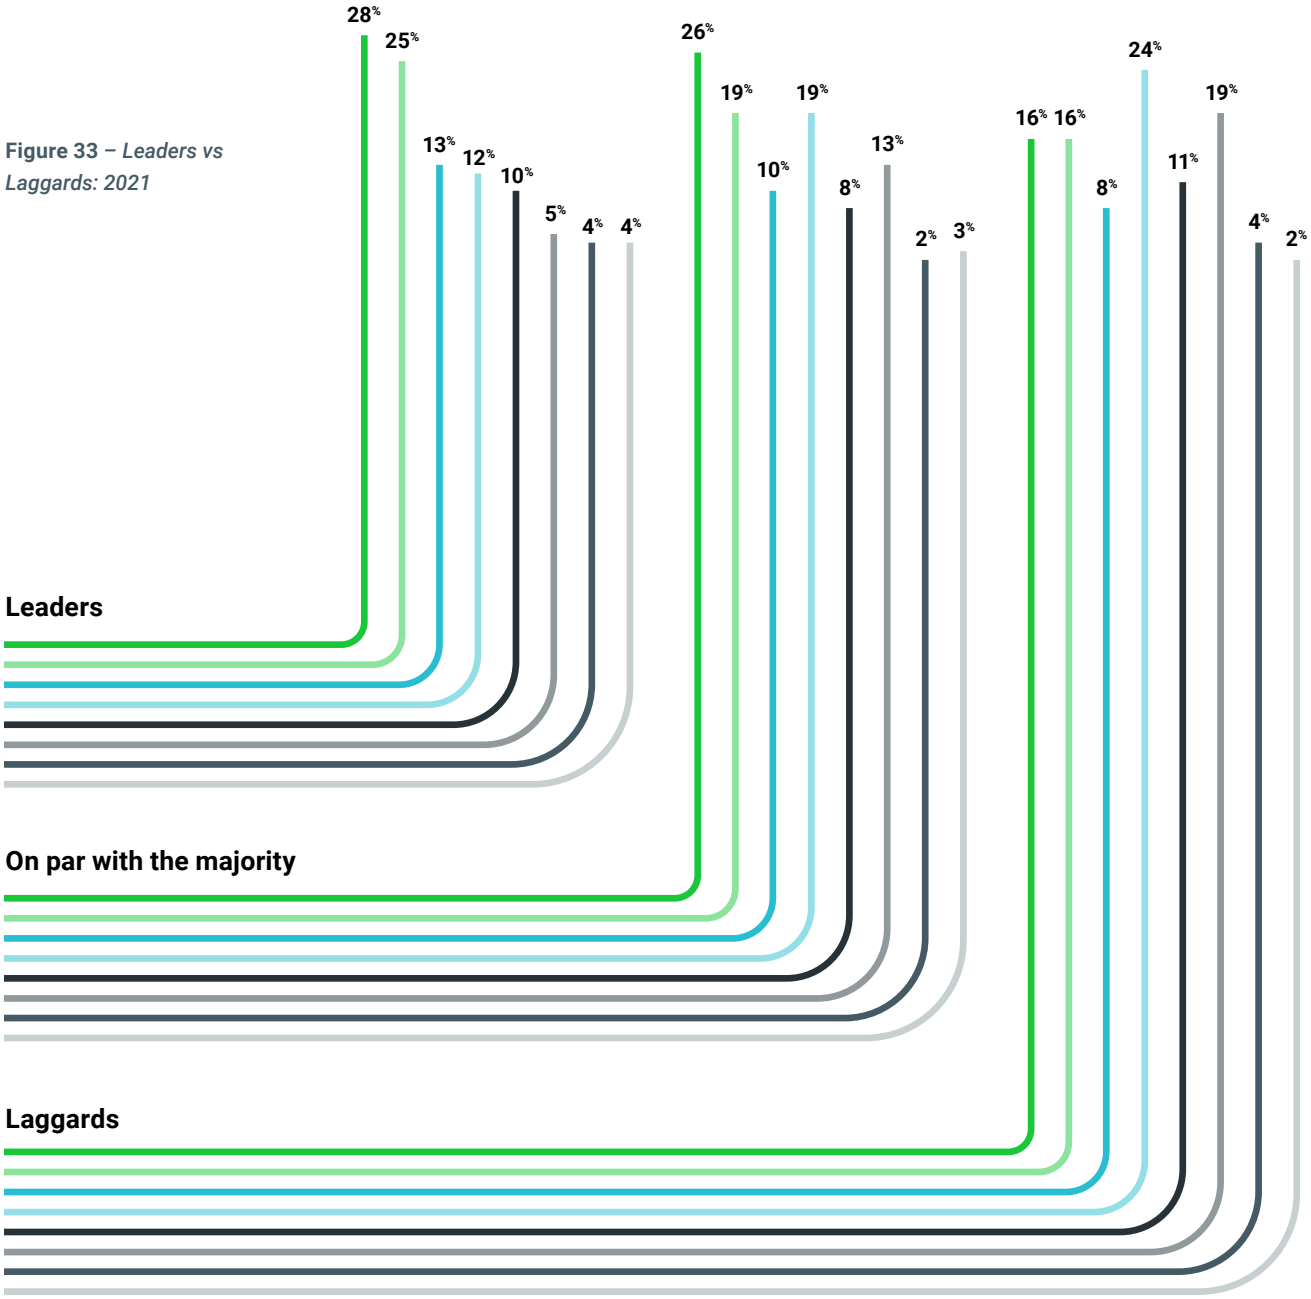

Optimize existing cloud resources for performance and cost

Advance a multicloud strategy (multiple public clouds with or without private clouds in the mix)

Modernize applications using public cloud containers and/or PaaS

Migrate more workloads to public cloud

Advance a (public) cloud-first strategy

We are not adopting public cloud

Automate governance/ compliance in public cloud

Implement O/CD in public cloud

In 2021, 67% of organizations are using Microsoft Azure, up from 61% in 2020

Last year, for the first time, Microsoft Azure overtook Amazon Web Services (AWS) as the leading cloud provider. They are continuing to maintain that lead. Also, notably, the percentage of respondents that are not using any form of cloud computing today dropped by half in 2021, from 8% to 4%.

Question: Which clouds are you using today? (Select all that apply)

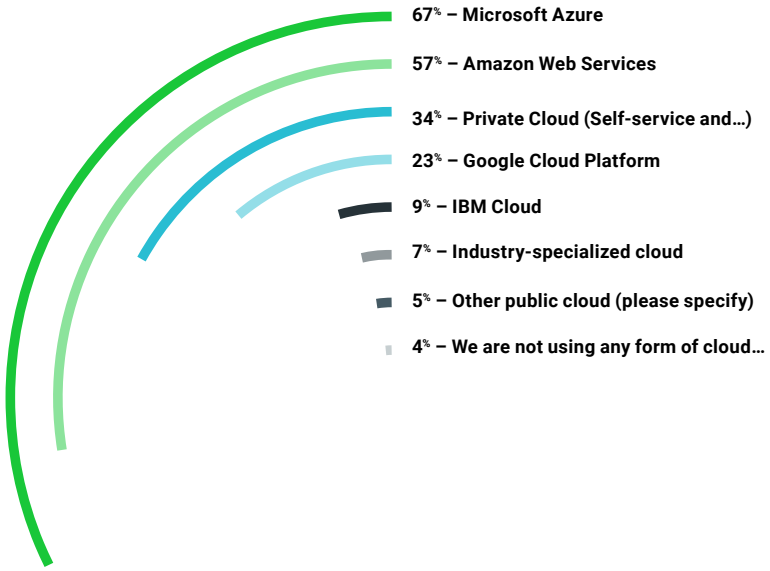

Figure 34

(N = 819)

30% of organizations are using 3 or more clouds today.

Multicloud is here to stay, but only 30% are using three or more clouds today.

The distribution of the number of clouds organizations are leveraging today.

Figure 35 – *Number of Clouds*

(2021: N = 819; 2020: N = 938; 2019: N = 736)

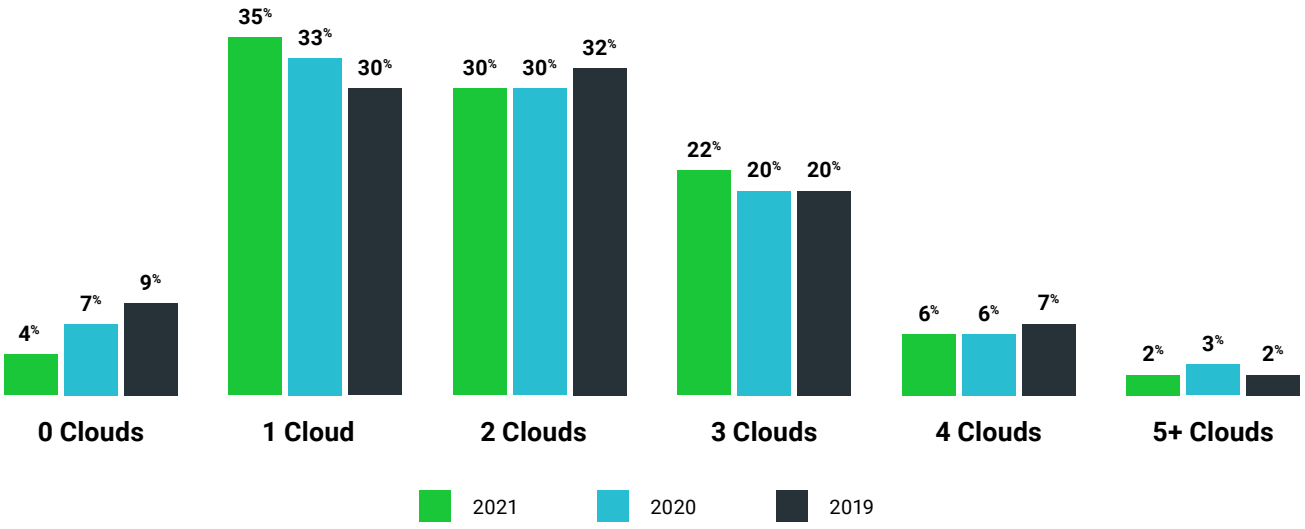

33% of organizations have a private cloud component

When we look at only those with a private cloud component, just 5% only have private cloud, while 8% are hybrid, and 20% are hybrid-multicloud (three or more clouds that include private cloud). Another way to put it, organizations with private clouds are more likely to be multicloud.

The distribution of the current number of clouds where at least one is a private cloud. Percentages are of total respondents.

Figure 36 – *Number of Clouds (with Private Cloud) (N = 278/819 (33%))*

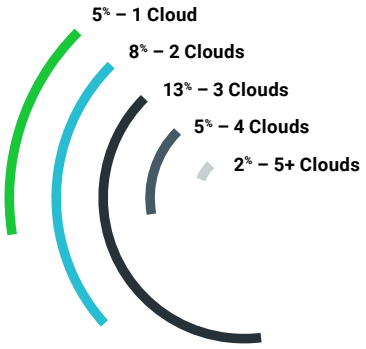

62% of organizations are using only public clouds with just 10% using 3 or more public clouds.

Looking at organizations that only use public clouds, about half (30%) are using a single public cloud, 22% are using 2 public clouds, and only 10% are using three or more public clouds.

The distribution of the current number of clouds where respondents are only running on public cloud(s). Percentages are of total respondents.

Figure 37 – *Number of Clouds (Public Clouds Only) (N=511/819 (62%))*

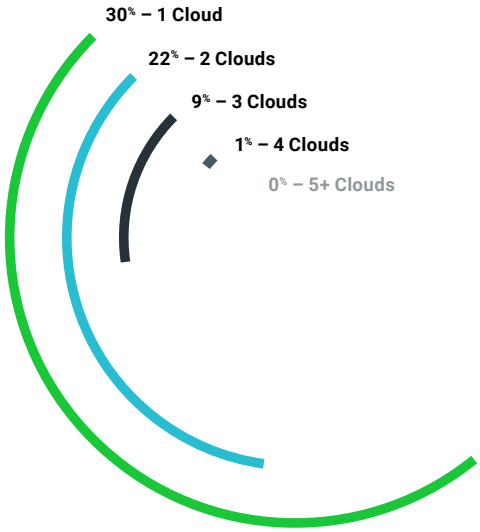

Azure and AWS will see the greatest investment in the next 12 months; other cloud providers must differentiate to compete.

Today, AWS and Azure are leading by a lot and many customers plan to spend more with them suggesting they will continue to outgrow the rest of the market.

More than half of respondents have no plans to adopt Google Cloud Platform (GCP) (55%), IBM Cloud (74%) or VMware (50%) clouds. Only a small percentage plan to increase spending in these clouds: GCP (16%), IBM Cloud (6%), and VMware on AWS/Azure/Google (18%).

These vendors have an uphill battle and must find ways to differentiate themselves, whether it is by the services they offer on top of IaaS or building on existing strategic partnerships with their customers.

When we look at how AWS, Azure, and GCP compare in terms of being used with additional clouds, the differences reveal some interesting patterns (**see fig. 39**).

Question: Over the next 12 months, does your organization plan to increase, decrease or maintain overall public cloud investments in the following?

Figure 38

N = 819

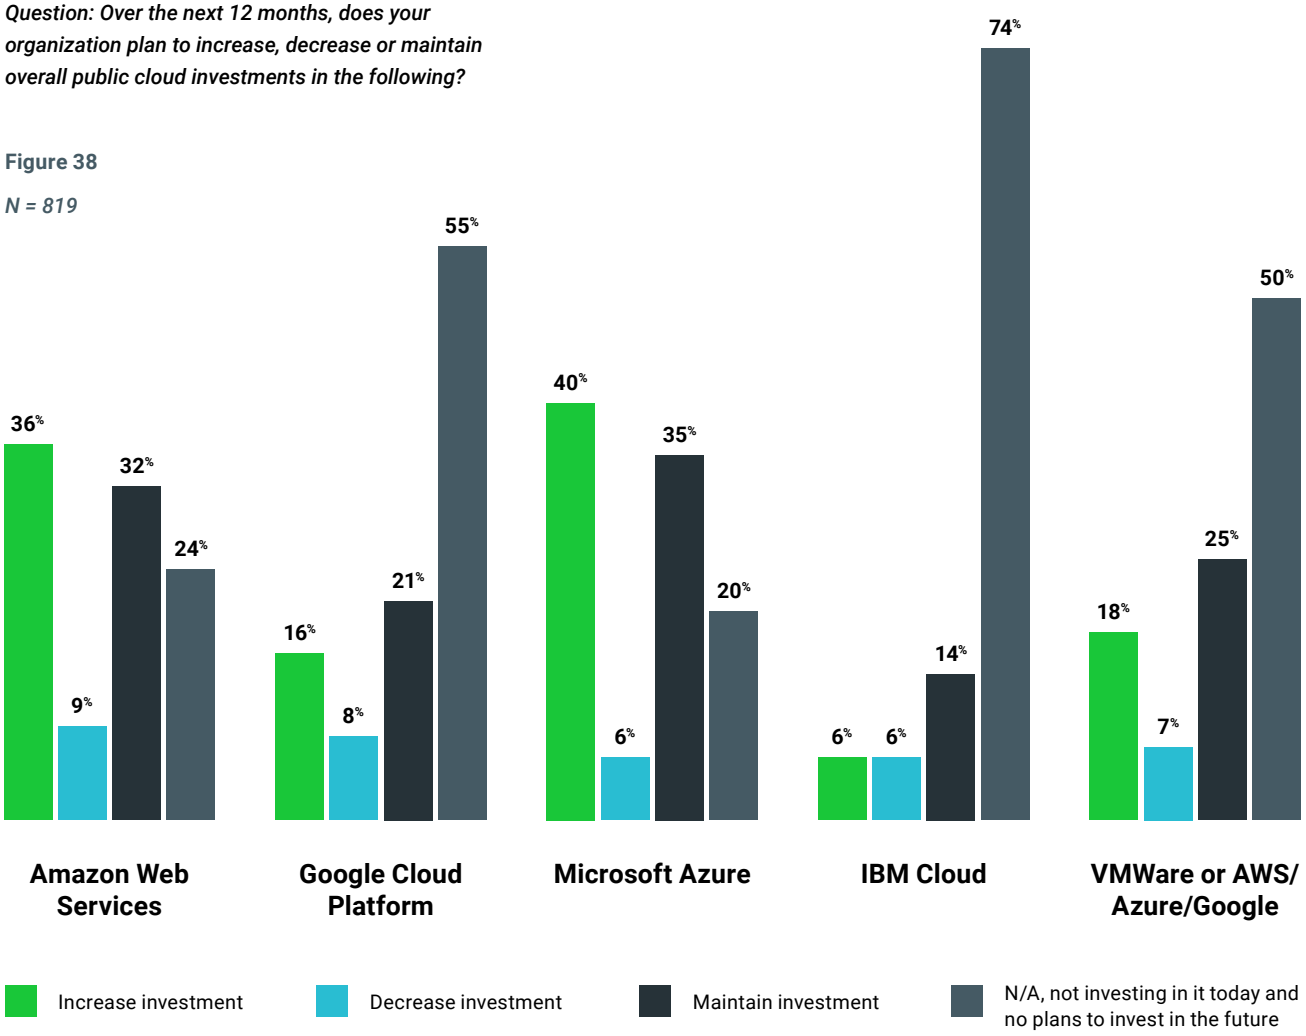

Google Cloud Platform (GCP) customers are much more likely to be leveraging the cloud provider as part of a broader mix of IaaS. It is indicative of it being late to the game—by the time GCP was an option, most companies had already adopted AWS and/or Azure. But, GCP is increasingly competitive on certain services, AI/ML services and Google Kubernetes Engine (GKE) in particular. For example, in December 2020 **Twitter announced it will be using AWS** as its main cloud IaaS provider (still running some elements off cloud). In **February 2021 they announced they will be using GCP** (they have used them since 2018) to "learn more from our data, move faster and serve more relevant content to the people who use our service every day." In other words, AI/ML.

For 23% of respondents using Azure, it is their only IaaS provider (this constitutes 15% of respondents as a whole). AWS is a distant second place at 13% (or 8% of respondents as whole). Azure's lead as a single public cloud provider is in part due to the fact that AWS competes in other industries, retail being an obvious one. Some organizations will prefer Azure as a result. Additionally, this lead could be the result of organizations having history as a Microsoft shop on-premises such that going cloud-first with a familiar partner was an obvious choice.

For each public cloud, how likely are they to be used alongside other public cloud providers. Percentages shown are the breakdown per cloud provider, ex. Azure's distribution is 23% + 33% + 31% + 9% + 4% = 100%.

Figure 39

Azure: N = 549; AWS: N = 469; GCP: N = 187

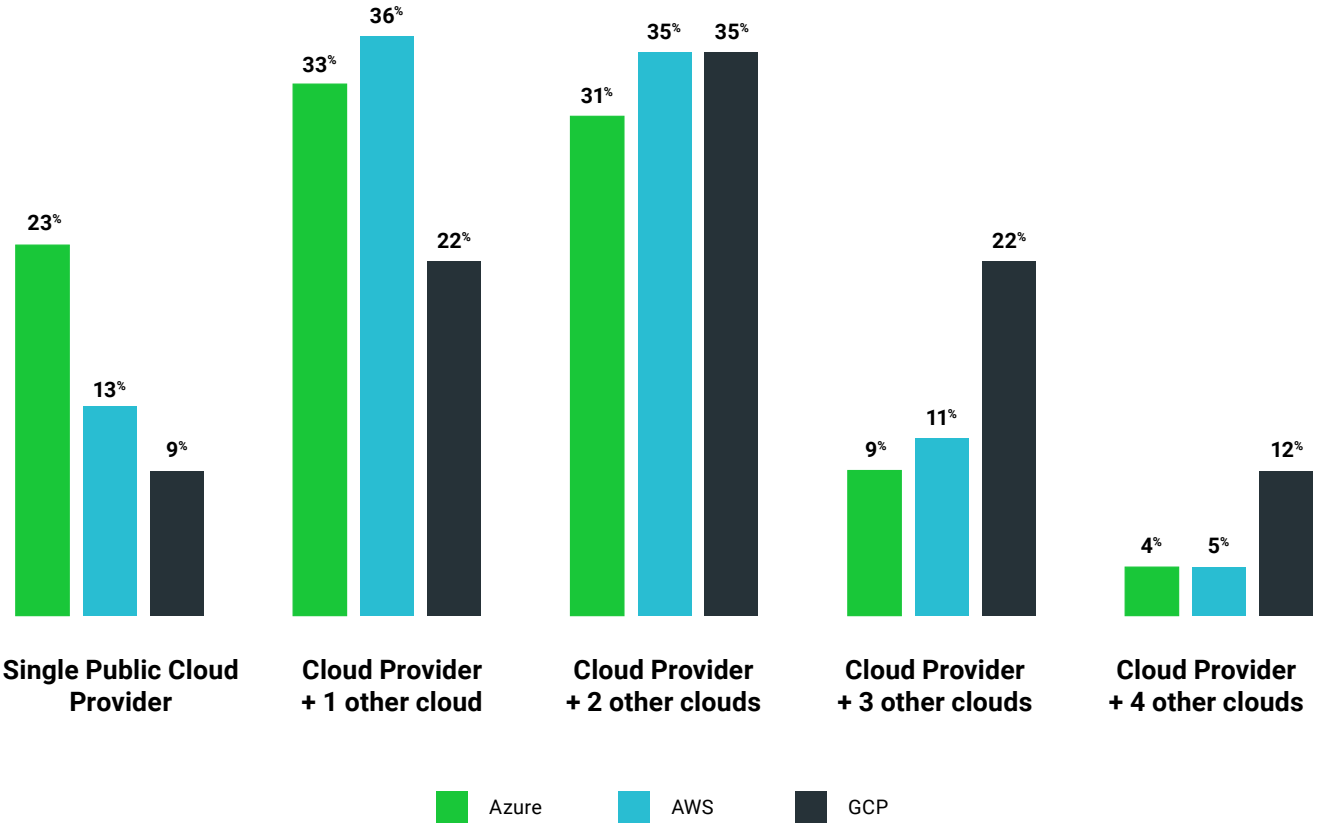

83% of organizations are using at least one cloud managed service today

We asked respondents to indicate what cloud managed services they are using today or planning to use within 18 months. Six hundred seventy-seven (677) of the 819 total respondents indicated current usage of at least one cloud managed service (83%). In 2020, 448 of 938 total respondents indicated current usage of at least one container platform add-on (48%).

Question: Is your organization making a strategic investment in the following types of cloud managed services? If not strategically investing in the service, please indicate that by skipping the row.

Figure 40

(2021: N = 752/819)

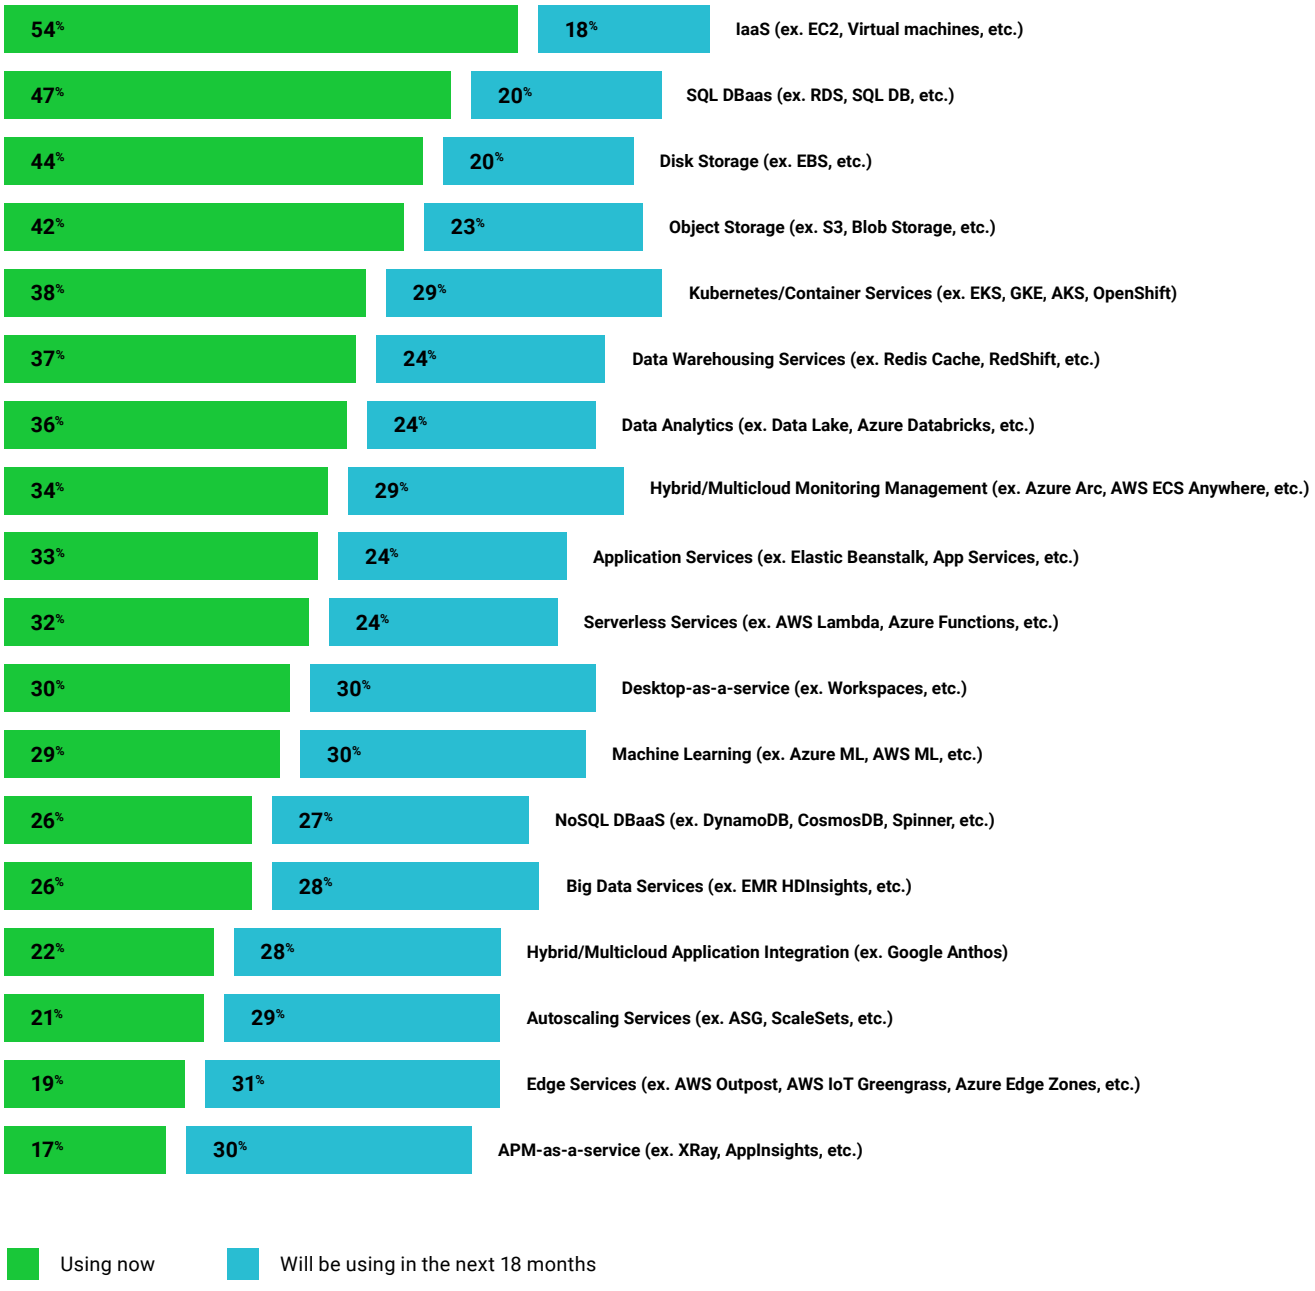

For 62% of organizations public cloud PaaS will play a strategic role for their business within 18 months.

Further indication of the importance of the service offerings on top of IaaS, is that most organizations expect that public cloud PaaS will play a strategic role for their business within 18 months. **Gartner anticipates that in 2021 PaaS adoption will grow at a higher margin than other public cloud services, by 26.6%**. They note,

"The increased consumption of PaaS is driven by the need for remote workers to have access to high performing, content-rich and scalable infrastructure to perform their duties, which largely comes in the form of modernized and cloud-native applications."

Question: By when do you expect public cloud PaaS solutions to play a strategic role for your organization?

Figure 41 – *Overall: 2021*

(N = 819)

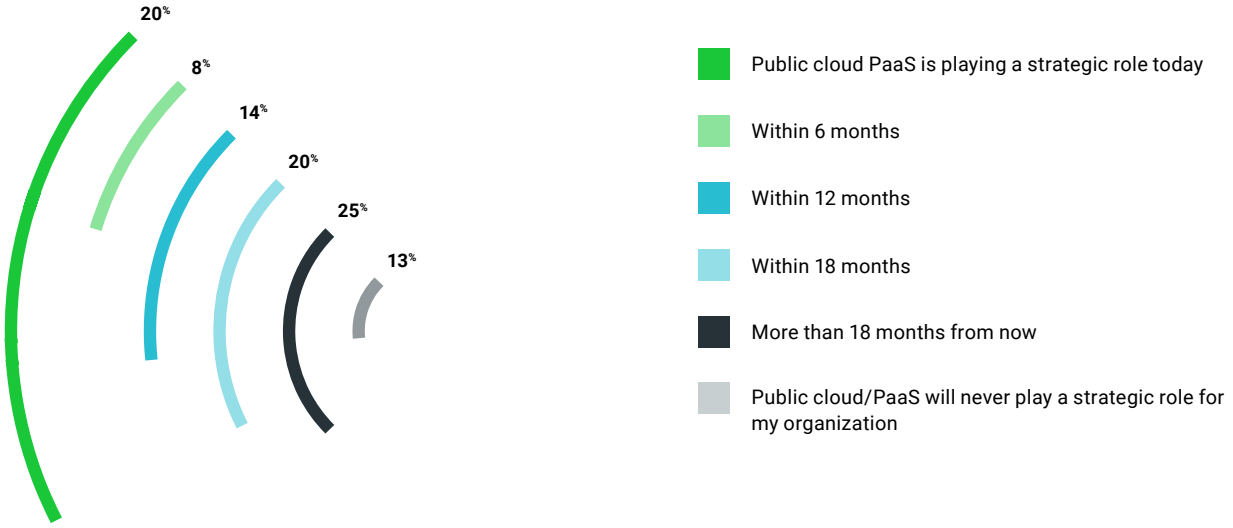

Figure 42 – *Leaders vs Laggards: 2021*

(Leaders - N = 279; Majority - N = 432; Laggards - N = 108)

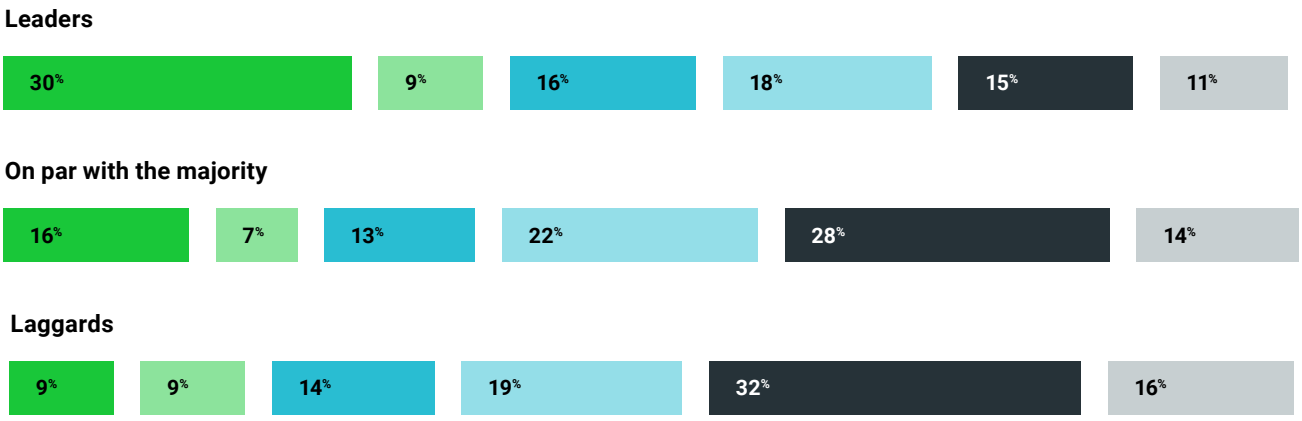

CHAPTER 4:

EDGE COMPUTING

Edge computing adoption is flat, but nearly 80% believe it is (or will be) relevant to their organization.

While edge computing only sees a slight bump in adoption, we will continue to keep an eye on this space. Seventy-nine (79%) believe that it is (or will be) relevant to their organization (**see fig. 45**). IDC anticipates that "by 2024, 25% of organizations will improve business agility by integrating edge data with applications built on cloud platforms, enabled by partnerships across cloud and communications service providers." [5]

Additionally, the Coronavirus pandemic has brought about a new normal for the future of work. IDC finds that "through 2023, reactions to changed workforce and operations practices during the pandemic will be the dominant accelerators for 80% of edge-driven investments and business model changes in most industries." [6]

Question: Is your organization leveraging edge computing today?

Figure 43 – *Overall: 2020-2021 (2021: N = 819; 2020: N = 938)*

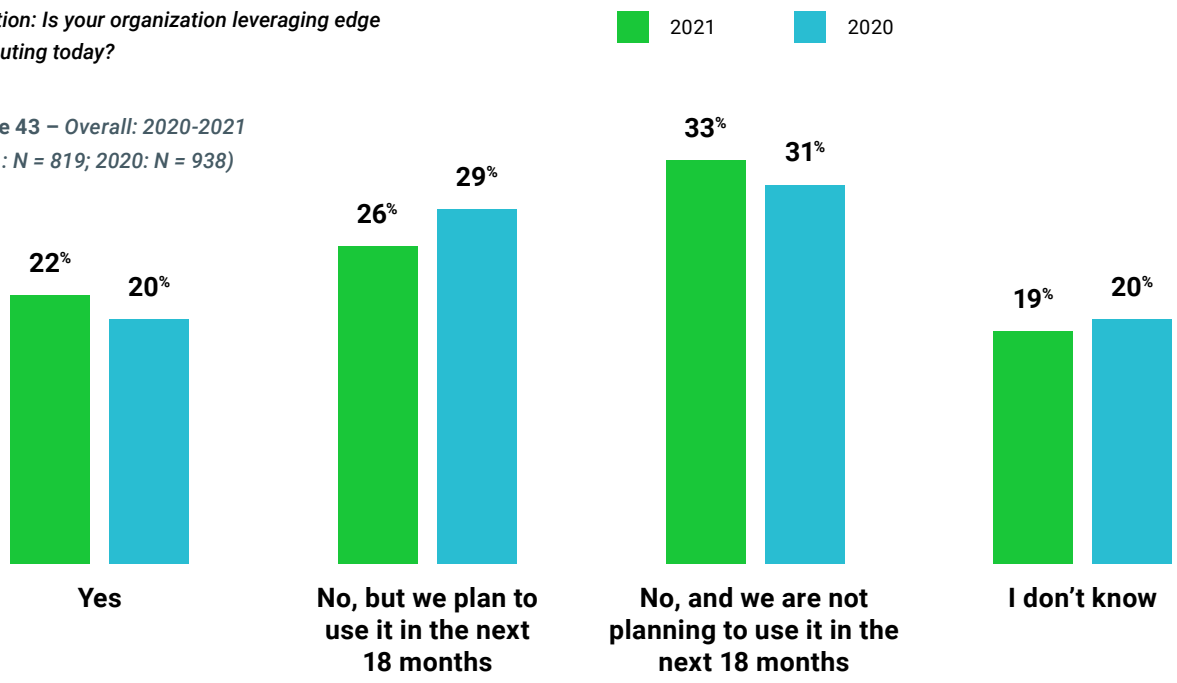

[5] IDC FutureScape: Worldwide Cloud 2021 Predictions. DOC #US46420120 / DEC 17, 2020

[6] IDC FutureScape: Worldwide IT Industry 2021 Predictions. DOC #US46942020 / OCT 27, 2020

Figure 44 – *Leaders vs Laggards: 2021*
(2021: Leaders - N = 279; Majority - N = 432; Laggards - N = 108)

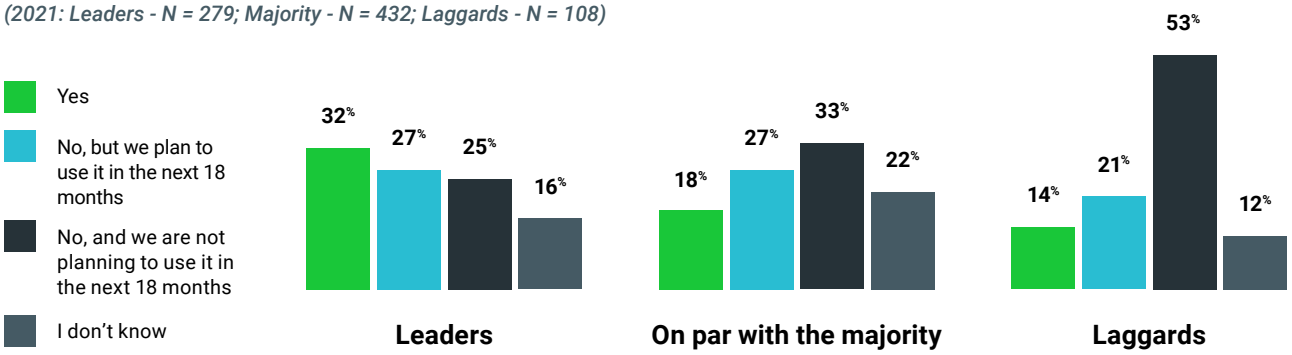

Minimizing latency is still considered the most relevant edge computing use case, with 32% believing it is (or will be) relevant to their organization.

Question: Regardless of whether you are doing edge computing today or not, which of the following use cases do you believe is (or will be) most relevant to your organization/business?

Figure 45 – *Overall: 2020-2021(2021: N = 819; 2020: N = 938)*

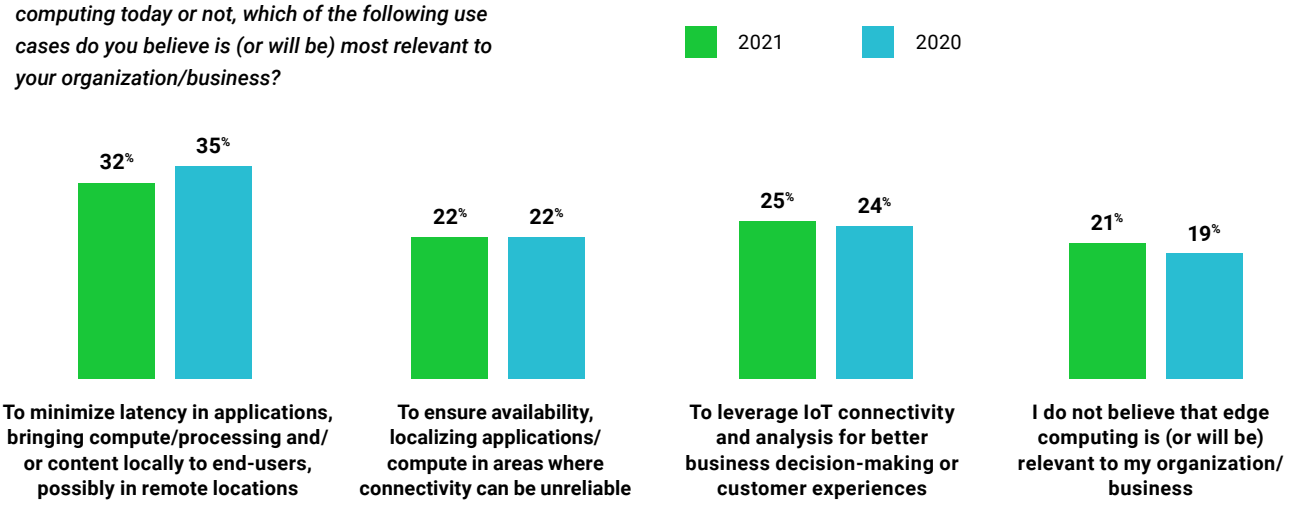

Figure 46 – *Leaders vs Laggards: 2021*
(2021: Leaders - N = 279; Majority - N = 432; Laggards - N = 108)

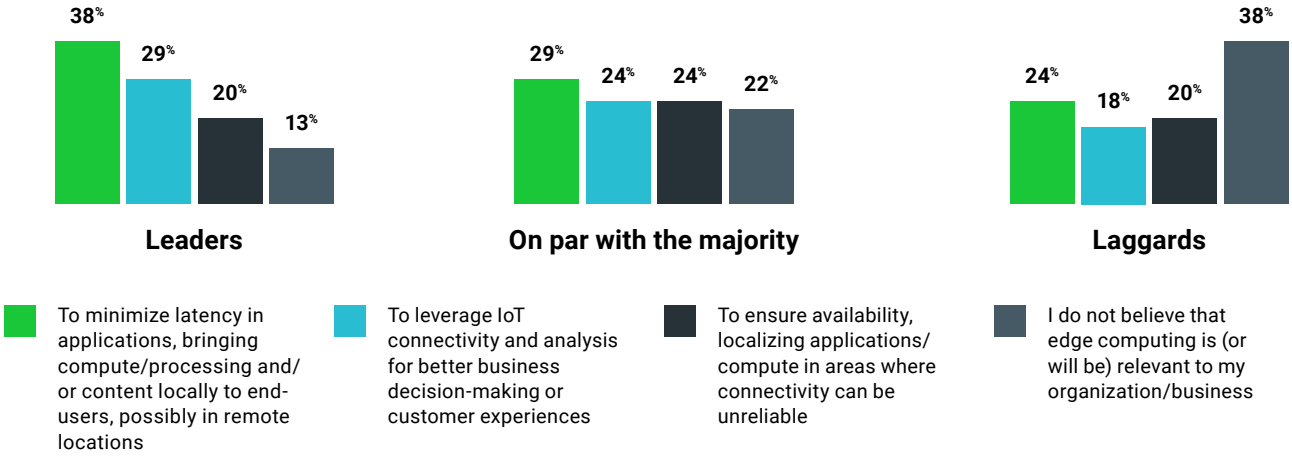

Complexity at 38% is the leading barrier by far to edge computing becoming conventional.

With edge computing there is an extraordinary amount of complexity by way of the distribution of data and processes and the tradeoffs that must be made. The laws of physics largely dictate that where the data lives, the application or process must run. Or you invite latency between the data and the application.

Figure 47 – *Overall: 2020-2021 (2021: N = 819; 2020: N = 938)*

Question: What do you believe is the primary barrier to edge computing becoming a conventional use case for organizations?

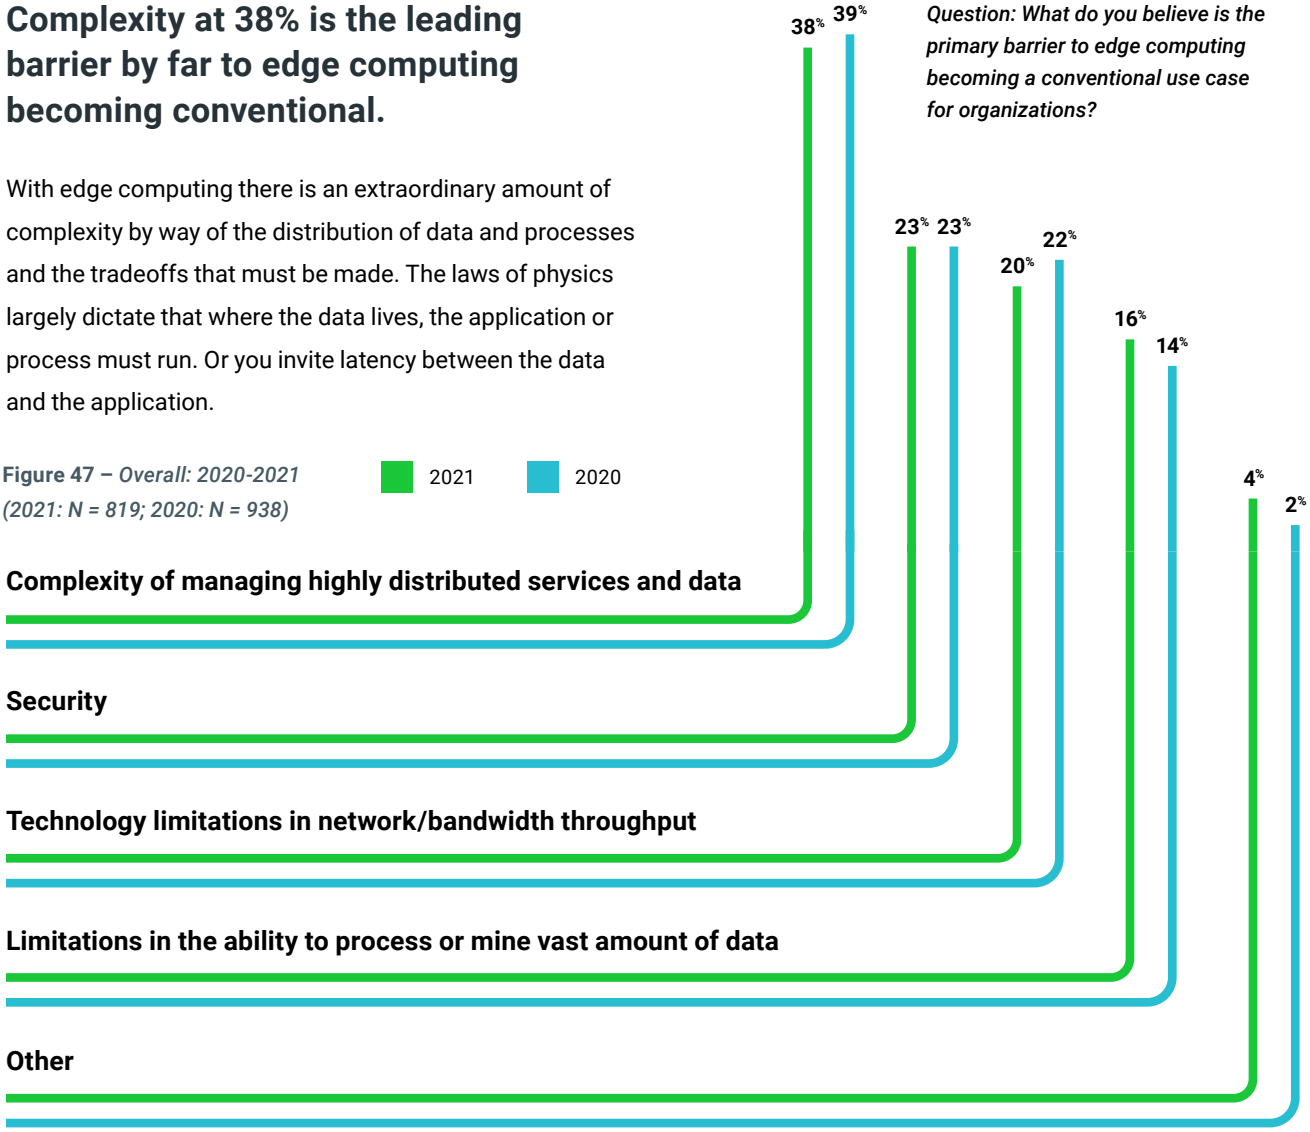

It is within these constraints that organizations must think about how they architect their applications for edge, as well as multicloud. However, a confluence of trends and technologies is raising its profile as a feasible business use case:

- Containers/Kubernetes enable the portability and orchestration of processes across clouds, local data centers, and local devices.
- Service mesh enables application services to communicate across this heterogeneous mix of infrastructure.
- 5G enables faster data speeds between the cloud and edge endpoints.
- And our ability to collect data at the source, created by people, places, and things, is rapidly improving.

The modern application will be a mesh of services operating and communicating across three main tiers:

- Endpoints such as sensors and IoT devices that collect data and perhaps execute some real-time processing, ex. a self-driving car stopping at a stop light.
- Local "edge" data centers that may run analysis and gather insights for consumption onsite, ex. a branch office or factory floor.
- Cloud with theoretically unlimited capacity where big data processing and AI/ML can be executed.

But tradeoffs exist across these tiers. Data can be processed faster locally, but there are limitations in compute and storage. The tradeoff is to have data processed elsewhere, but introduce network latency in getting that data to where it needs to be. You have to bring the data through these tiers and the question becomes what processes can you move where? How much can (or "must," think self-driving cars...) be processed in real-time at the end-point? How much can be processed locally where you get the benefits of minimizing latency in data processing, but have limited compute and storage capacity? Alternatively, you can process that data in the cloud with its theoretically infinite capacity, but you take a hit on network latency, as determined by the capacity of the network. Ultimately, it's a tradeoff of limited compute and storage capacity at the edge vs. network delay that is the result of moving that data to be processed elsewhere.

These resource tradeoffs have existed since the dawn of virtualization. The complexity exists today and will only increase. And, as our survey finds, that complexity is expected to be the primary barrier to edge computing becoming a conventional use case.

66% are optimistic about edge computing's market offerings.

Question: What is your perception of the market's current offerings (open source and commercial) that would allow your organization to leverage edge computing?

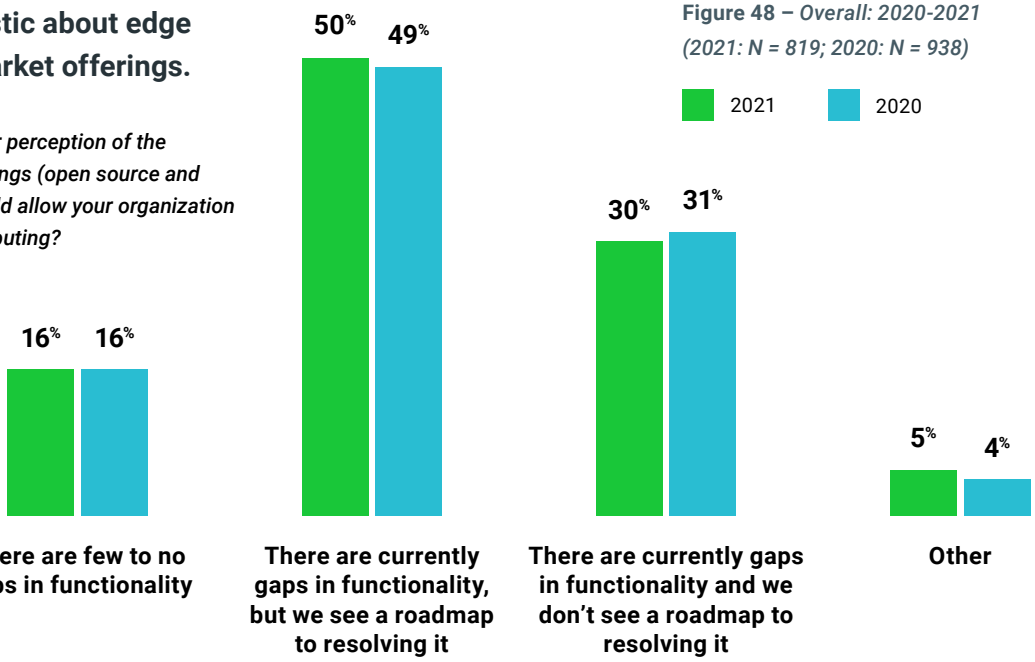

Figure 48 – *Overall: 2020-2021 (2021: N = 819; 2020: N = 938)*

ABOUT TURBONOMIC, AN IBM COMPANY

Turbonomic, an IBM Company, provides Application Resource Management (ARM) software used by customers to assure application performance* and governance by dynamically resourcing applications across hybrid and multicloud environments. Turbonomic Network Performance Management (NPM) provides modern monitoring and analytics solutions to help assure continuous network performance at scale across multivendor networks for enterprises, carriers and managed services providers.

For further information, please visit www.turbonomic.com

*www.turbonomic.com/resources/case-studies