

# Kubernetes and cloud native operations report 2022

Data from 1300 respondents on hybrid and multi-cloud operations, Kubernetes, VMs, bare metal, goals, benefits, challenges, operators, advanced usage, edge, and more.

## Introduction

Canonical, the makers of Ubuntu, is a member of the Cloud Native Computing Foundation ([CNCF](#)) and provides commercial support for various technologies that are part of the CNCF ecosystem. As an active part of the community, we take the pulse, analyse and share insights on personal and commercial use of cloud native technologies, leveraging the Ubuntu user base and our proven experience working with open-source software and solving enterprise complexity. We contribute the data back to the community, along with our own analysis and the insights of industry experts in the form of commentary. The goal is to help improve the cloud native technologies to best address the needs of end users and their organisations.

## Key takeaways

- **Hybrid vs multi-cloud: the reality behind the adoption**

More than 83% of respondents told us they're using either hybrid or multi-cloud. In the last year alone, the percentage of respondents who did not use hybrid or multi-cloud dropped from 22.4% to 16.4%. This time, Tim Hockin discusses the reality behind that adoption: "People often build a straw man of hybrid or multi-cloud, with the idea of one giant mesh that spans the world and all the clouds, applications running wherever capacity is cheap and available. But in reality, that's not at all what people are doing with it. What they're actually doing is using each environment for just the things they have to use it for."

- **Kubernetes on bare metal**

The question of where to run applications is an interesting one. 14% of respondents said that they run everything on Kubernetes, over 20% said on bare metal and VMs, and over 29% said a combination of bare metal, VMs, and Kubernetes. As highlighted by Ihor Dvoretzkyi, this distribution shows how the flexibility of Kubernetes allows organisations to run the same type of workloads everywhere. Looking back at last year's highlight, where Kelsey Hightower stated that bare metal was a better choice for compute and resource-heavy use cases such as interactive machine learning jobs, it seems that the tune is changing. Actually, as running Kubernetes is becoming more accessible, Alexis Richardson speculates that [organisations would further adopt Kubernetes on bare metal if they knew it was possible](#).

- **Security is still everyone's concern**

This comes as no surprise: 38% of respondents suggest that security is the most important consideration whether when operating Kubernetes, building container images or defining an edge strategy. Keeping clusters up-to-date is a definitive best practice to solve security issues. However, according to Jose Miguel Parrella, it is not as embedded within the strategic-thinking IT infrastructure group as one could expect. Today, it is more of a Day-30 discussion that only occurs within the small team of Kubernetes maintainers of every organisation. Combined with the fact that only 13.5% of people reported that they've "mastered" security in the cloud native space, it is clear that organisations have some room to grow when it comes to properly adopting and managing Kubernetes in production.

- **An app store for operators**

When asked if they would trust an operator built by an expert, more than 50% of our respondents said yes. This makes sense considering that the skills gap is a major issue for organisations. However, the provenance and accessibility of operators need to be addressed to mitigate the main concerns of organisations adopting new technologies, and in particular open-source solutions. As the automation of operations continues to grow, finding a safe place to get the necessary tools is becoming more and more important. In both reports, Karthikeyan Shanmugam and Alexis Richardson spoke about the idea of an "App Store" where people can publish and consume operators, similar to [Charmhub](#).

- **What is in store for the future?**

Despite the obstacles, Kubernetes adoption is consistently growing. It is only fair to try and think about what the future holds. For Ihor Dvoretskyi, the high-level goals — improving maintenance, monitoring and automation, and infrastructure modernisation — are likely to stay the same in the years to come. However, use-case related evolutions are to expect, in particular in the AI/ML and data platforms space. And as the platform evolves to support more diverse tools, the goals of its users will evolve with it.

## About the survey

The Kubernetes and cloud native operations survey used to collect the data for this report ran in late November 2021 after Kubecon North America. Nearly 1300 people responded to the survey. We collected expert 3rd party commentary on the results from the public clouds, ISVs, and from an expert in the financial industry who also co-chairs the Operator Framework at the CNCF. We're happy to share this information with you and look forward to running the survey again, later in the year!

# Kubernetes and cloud native operations report 2022

## Who is using Kubernetes and cloud native technologies?

### 1. Respondents by job role

I am best described as:

1216 out of 1279 people answered this question

16.9%	Infrastructure Architect	205 responses
10.4%	Site Reliability Engineer	127 responses
9.6%	Full-Stack developer	117 responses
9.2%	Platform Engineer	112 responses
9.1%	IT Manager	111 responses
8.3%	Application Developer	101 responses
7.6%	Academic/Teacher/Student	93 responses
6.4%	Software Architect	78 responses
5.6%	Consultant	68 responses
5.4%	Back-end Developer	66 responses
4.9%	Security Engineer	59 responses
2.9%	Business Executive	35 responses
2.8%	Data Scientist	34 responses
0.8%	Other	10 responses

The majority of Kubernetes users are people that are building platforms or interacting with platforms, as Kelsey Hightower noted last year, and this year is no different. They might wear several hats, but all cloud native users aim to add value to their business by leveraging those technologies. Here's what our experts had to say:

"There are differences between the words you use to describe your job, versus what your job actually is. This will be one of the areas that will change the most. The skills will not change dramatically, but the jobs that people do will evolve. And the way the people describe those jobs will change too."

—Alexis Richardson, Founder and CEO, Weaveworks

“Clearly, the people who are most interested in Kubernetes are the people who have to operate software across clusters. K8s is a cluster operations API. It allows us to think about a cluster as a unified thing, and then describe what we want on that cluster. And it takes care of the details. So it makes a lot of sense to me that architects and DevOps/SREs are the ones who topped the charts here. Those are groups that think about how to operate software cleanly in an organisation, across multiple clouds. Can this be done? Kubernetes gets us a lot closer to saying “Yes” on a question like that.”

—Mark Shuttleworth, Founder and CEO, Canonical

## 2. Respondents by industry

My company is in the following industry:

1243 out of 1279 people answered this survey

36.9%	Software/Technology	459 responses
10%	Education	124 responses
7.4%	Financial Services	92 responses
7.4%	Consulting	92 responses
6.8%	Telecommunications	80 responses
3.8%	Professional services	47 responses
3.7%	Healthcare and life sciences	46 responses
3.5%	Manufacturing	43 responses
3.3%	Government	41 responses
2.9%	Retail and e-commerce	30 responses
2.2%	Consumer	27 responses
2.2%	Scientific or technical services	27 responses
1.9%	Construction	23 responses
1.7%	Media	21 responses
1.4%	Energy and Utilities	17 responses
1.2%	Agriculture	15 responses
1%	Transportation and warehousing	13 responses

One of the positive surprises from the [2021 report](#) was the ranking of financial services in the second place of the most represented industry, as noted by Ken Sipe and Kelsey Hightower. FinServ is adopting Kubernetes while navigating the security and compliance implication of its integration. This year, Education is the lucky second, with the Software/technology industry still holding the reins. Here's what Tim Hockin had to say:

“This gels with what I hear from people, they're spending an enormous amount of time and energy on IT, which is that whole first category. IT is generally a cost centre. Folks would like to spend less money on that. Modernizing is part of it, but as a goal on its own this very much lines up with what we see in the industry. Typically modernising means adopting containers and cloud native solutions, switching from bespoke technologies to CNCF tools; switching from legacy monitoring systems to Prometheus; or switching from vSphere to Kubernetes.”

—Tim Hockin, Principal Software Engineer, GCP

# Cloud native: goals, benefits, and estate size

## 3. Which technology goals are most important?

Which goals are most important to you and your team? Choose 2.

1276 out of 1279 people answered this question (with multiple choice)

64.3%	Improved maintenance, monitoring, and automation	821 responses
43.9%	Modernizing infrastructure	560 responses
25.9%	Faster time to market	331 responses
17.8%	Lower infrastructure TCO	227 responses
15.1%	Removing vendor dependencies	193 responses
11.6%	Global reach	148 responses
10.4%	Agility around traffic spikes	133 responses
9.5%	Ensure portability	121 responses
1.3%	Other	17 responses

Karthikeyan Shanmugam rightfully highlighted last year that the top results were consistent with the typical concerns of the biggest share of the audience of the survey. Actually, the top reasons cited are very representative of SRE's concerns: 77.8% of SREs/DevOps selected improved maintenance, monitoring and automation as one of their two choices for this question. It is consistent with what we are seeing this year. Automation keeps coming back as a big theme, as it helps to lower operating costs and ties into faster time to market, which is definitely a good business driver.

"I do not expect the high-level goals to change too much in the years to come. People are sticking to their habits and their top priority will mostly be about running software on top of the reliable, stable infrastructure—they actually do not want to have to care about the infrastructure. It reminds me of the old joke that says "A good sysadmin doesn't do anything because they have set everything properly. A bad sysadmin has to work a lot as they probably have to fix something broken all the time.". If I think of this question with regards to problem spaces and tools I can definitely say that I expect this to change. Two years ago, the ecosystem was taking the most interest in serverless technologies and that is still the case today. But today I also see a lot of emerging interest in AI/ML and data platforms on Kubernetes. This is mostly because of the platform's growing maturity. Kubernetes started as a platform for stateless applications but has grown to support stateful too. So I can imagine that as the platform evolves to support more diverse tools the goals of its users will evolve with it."

—Ihor Dvoretzkyi, Senior Developer Advocate

"If you are just getting started with cloud native technologies, I'd advise you to double down on automation and modernization: this has always been the most important benefit. If you are starting your cloud native journey thinking about removing vendor dependencies, you're missing the opportunity of improved feature velocity."

—Jose Miguel Parrella, Principal Program Manager, Office of the Azure CTO

## 4. What are the top benefits of cloud native technologies for businesses?

Kubernetes and cloud native technologies unlock innovation for organisations and allow them to achieve their goals. The benefits of cloud native technologies vary, depending on their usage and the maturity of the organisations using them. Last year, elasticity and agility, resource optimisation and developer productivity occupied the top spots amongst our users, regardless of their job group. This year, the results are different and our experts' opinions also highlight how diverse those benefits really are.

Which goals are following do you consider the most significant cloud-native benefits? Choose 2.

1271 out of 1279 people answered this question (with multiple choice)

50.3%	Elasticity and agility	639 responses
26.5%	Resource optimization	337 responses
26.5%	Resource optimization	337 responses
21.4%	Reduced service costs	272 responses
21.2%	Faster time-to-market	270 responses
19.4%	Cloud portability	247 responses
18.6%	Developer productivity	236 responses
15.1%	Open-source software	192 responses
13.5%	Simpler operations	171 responses
8.7%	Cutting-edge technology	110 responses
5.4%	Global reach	68 responses

"If you are doing open source, you should look at it just like as if you were buying software. When you buy software, you're paying an institution to have the highly technical people on staff. When you do open source, you might need to hire those people yourself or pay a company like Canonical to do that for you.

Cost of Ownership is not to be misconstrued with the maintenance costs that every software should and does have. It really is less about lower infrastructural costs and more about lowering risk. I can take the risk completely off the table because I have the software and if I need to, I can totally rebuild what I need."

—Ken Sipe, Cloud Native Computing Foundation, Edward Jones

"People really benefit from a faster time to market and agility, but it's also interesting to note how the composable nature of containers has introduced open source to a whole new set of people and organisations."

—Jose Miguel Parrella, Principal Program Manager, Office of the Azure CTO

“I’m a big believer in “don’t fix things that aren’t broken”. So if you look at microservices as a panacea, then you’re going to be disappointed. Rather, if you look at it as a means to solving a specific set of problems that many folks have—but not everybody—then you will be better off. It’s a way of organising teams and of empowering Conway’s Law in some sense. It’s a law for a reason. Microservices provide a good way of doing that. But if you think it’s going to take a bad application and make it good, then you’re going to be disappointed. Or if your application is unreliable, or it follows the big ball of mud architecture, then you’re also going to have a hard time.”

—Tim Hockin, Principal Software Engineer, GCP

## 5. What is the definition of “hybrid cloud”

How do you define a “hybrid cloud”

1232 out of 1279 people answered this question

79.1%	A combination of at least one private cloud and one public cloud	974 responses
10.8%	A combination of at least two public clouds and a private cloud	133 responses
4.4%	A combination of at least two public clouds	54 responses
4.2%	A combination of at least two private clouds	52 responses
1.5%	Other	19 responses

We see more and more people getting the definition of hybrid cloud right. Mark and Ken provide a forward-looking perspective, now that this has become common knowledge:

“The key question, really, is how much of what you do every day can you do on multiple different clouds without thinking about it? For me, the sensible thing for a medium or large institution is to have a fully automated private cloud and also relationships with at least two public cloud providers. This way, businesses essentially benchmark themselves on doing any given operation on the private cloud and on the two public clouds. Not many organisations are there though. I hear people giving very shallow reasons for being stuck on a particular cloud. Good practice dictates that you have a pretty high bar before you lock a particular workload to a particular cloud and make sure you explore and exploit the unique characteristics and services of each cloud.”

—Mark Shuttleworth, Founder and CEO, Canonical

“The mentality that there’s a difference is legacy thinking: the cloud is just a data centre. It’s just a matter of how much control you have over it, and what you’re willing to pay for.”

—Ken Sipe, Cloud Native Computing Foundation, Edward Jones

## 6. What are your hybrid cloud or multi-cloud use cases?

We mainly use hybrid or multi-cloud for/to:

1228 out of 1279 people answered this question

22.1%	Accelerate development, automate DevOps	272 responses
16.4%	We don't use hybrid or multi-cloud	201 responses
15.1%	Disaster recovery	186 responses
15.1%	Expand cloud backup options to cut costs	185 responses
7.4%	Move an application	91 responses
7.3%	Cluster mission-critical databases	90 responses
5.9%	Switch between public cloud providers in a flash	73 responses
5.1%	Cloud bursting	63 responses
3.3%	On and off-ramp data	40 responses
1.3%	Monitor and predict usage costs	16 responses
0.9%	Other	11 responses

Compared to the 2021 report, more respondents reported using hybrid/multi-cloud technologies to accelerate development and automate DevOps. Where the attitude towards hybrid/multi-cloud technologies leaned more towards cautious trials, it seems that the consensus is progressively shifting into considering possible optimisations based on the strengths of the different platforms.

Relevant Quotes:

“Given the size of the application landscape in enterprises, understanding the incremental changes to the enterprise environment as a result of moving workloads to the clouds is much needed in a hybrid or multi-cloud scenario.”

—Alexis Richardson, Founder and CEO, Weaveworks

“Cloud has done a wonderful thing, which is to make technology, and in particular, operations, an economic question; “Where can I do this most effectively?” In the old days, there was no choice, there was a monopoly, and that was everyone’s internal IT. That monopoly is not broken yet, but at least internal IT now has some competition. As IT you compete on a CapEx basis with OpEx. So you have to go about it slightly differently and be able to report savings that make sense to the CFO.”

—Mark Shuttleworth, Founder and CEO, Canonical

“People often build a straw man of hybrid or multi-cloud, with the idea of one giant mesh that spans the world and all the clouds, applications running wherever capacity is cheap and available. But in reality, that’s not at all what people are doing with it. What they’re actually doing is using each environment for just the things they have to use it for. For example, they may have a giant footprint on AWS, and have a gravity well of data there—so they’re run some stuff on AWS, and they want to use BigQuery on GCP so they run some stuff on GCP. Or they have things that they just aren’t willing or able to move to the cloud, and so they keep it on-premise. But these things are largely disconnected from each other, or the connections between them are very carefully controlled.



There are also a lot of enterprises, like FinTech, that are still conscious of security or sovereignty. And they say, “I know this works, my auditors are happy. So I’m going to keep my critical things in my own closet, and I can move all the satellite things off. But the gravity well is in my data centre.”

—Tim Hockin, Principal Software Engineer, GCP

“A lot of organisations will use hybrid to be closer to their customers, to respond faster to changes. For example, a change in the Yocto world requires you to go on location with a flash drive to update your system. Even if not everyone is currently using hybrid-cloud, there’s no denying that there’s a trend: every cloud provider is putting Kubernetes wherever their customers are, closer to developers.”

—Jose Miguel Parrella, Principal Program Manager, Office of the Azure CTO

## 7. Size is important: How many machines and clusters are people running?


On average, how many machines are in your fleet (incl VM, bare metal, etc)?

1167 out of 1279 people answered this question



If you use Kubernetes, how many production clusters do you have?

1242 out of 1279 people answered this question



The size and number of clusters is dependent on multiple factors, from the company size and security requirements to the familiarity with best-practice deployment. It is an interesting indicator of how a company is thinking about their platforms. Compared to last year, we see a 2.2% increase in respondents that manage more than 500 machines and steady numbers of Kubernetes production clusters. Our experts had great insights to share on this:

“The number of clusters will exponentially grow and more of them will be transient. The order of magnitude of clusters will be in the 10.000, as many clusters as there are people in an organisation.

People will create their own clusters and throw them away all the time for all kinds of development and CI/CD. An enterprise with 5000 apps, would run those on about 1 to 5000 clusters. Why not, as these can just be disposable. A team responsible for a piece of functionality, say five apps, could run those five apps in one or two big production clusters. But they might also have multiple regions. I wouldn't be surprised to see as many clusters as apps, but that doesn't imply a 1 to 1 mapping.”

—Alexis Richardson, Founder and CEO, Weaveworks

---

“Regarding Kubernetes cluster size, I don't see the numbers dramatically increasing here. This makes sense - you ran the last survey 6 months ago and people are being cautious about moving their production workloads to Kubernetes. That aligns with the skills gap data point. I do expect more businesses will have Kubernetes clusters in production in the years to come though, and as you can see, the number of people that do not have any K8s production cluster has also slightly decreased.”

—Ihor Dvoretzkyi, Senior Developer Advocate

---

“All the work that is going on in the ecosystem, with initiatives such as multi-cluster, federation, service mesh and fleet management, will eventually blur the boundaries between clusters. I believe that the different approaches regarding cluster sizes and developer workspaces, whether that's a dedicated cluster or dedicated namespaces will eventually converge. In four or five years' time, it will be easy to interact and share resources between namespaces in a Kubernetes cluster, and across clusters. And that's one argument less for this single, giant cluster, to exist.”

—Michael Hausenbla

---

# Container and Kubernetes usage

## 8. Applications platforms

Where do you run applications in your organisation?  
Choose the most accurate statement.

1229 out of 1279 people answered this question

29.1%	On a mix of bare metal, VMs and Kubernetes	358 responses
20.3%	On a mix of bare metal, VMs	250 responses
15.3%	On VMs, evaluating Kubernetes for development	188 responses
14.6%	Mostly on VMs, planning a full migration to Kubernetes	180 responses
14%	All our applications run on Kubernetes	172 responses
5.6%	Don't know	69 responses
1.0%	Other	12 responses

The choice of substrate, or the base infrastructure layer used, is closely linked with the type of workloads and the tools used to manage them. As Kelsey Hightower shared in last year's report, Kubernetes is a powerful tool, but bare metal is a better choice for compute and resource-heavy use cases such as interactive machine learning jobs. However, it is getting easier to run Kubernetes on bare metal. These considerations are reflected in our experts' answers:

"I suspect that the unfortunate truth is that we all look at the potential market for cloud and we say it's, you know, xx billion dollars on-prem, and everybody wants to move to the cloud, but the movement is really slow. And three years from now, we may see the top line move up to 35% or 40%. But even 40% feels aggressive—we are below 29% now, so I guess it'll probably be in the low 30s in three years. I guess people running on VMs today planning a full migration to Kubernetes over a three-year timeframe will be a significant driver of that change."

—Tim Hockin, Principal Software Engineer, GCP

"This highlights the fact that people are trying to mix the different types of infrastructure for solving their different needs. In this case, the flexibility of Kubernetes is a significant benefit. If people are using Kubernetes they can run it basically everywhere. On a laptop, on a bare metal machine, in a private data centre, or on a public cloud VM. As long as people are using Kubernetes, as a necessary slice at the top of the infrastructure layer, right below their applications, they are using a universal set of tools and APIs. This allows them to run the same type of workloads everywhere."

—Ihor Dvoretzkyi, Senior Developer Advocate

“People using bare metal would probably be more interested in Kubernetes if they knew it could run on bare metal. Running K8s on top of bare metal is quite difficult and new. They’re using bare metal for a reason, so they’ve ruled out using vSphere. They probably think of K8s as being a bit like that and so they’re not sure if they want to use it.”

—Alexis Richardson, Founder and CEO, Weaveworks

## 9. What cloud native use cases are you working on?

Which of these cloud-native use cases are you working on now?

1258 out of 1279 people answered this question

19.1%	Re-architecting proprietary solution into microservices	240 responses
15%	Deploying and testing applications in a CI/CD pipeline	189 responses
13.4%	Moving to an open-source solution	168 responses
10.7%	Managing or enabling a hybrid-cloud setup	134 responses
10.1%	Deploying or managing Kubernetes-as-a-Service	127 responses
9.9%	Orchestrating workloads across a multi-cloud setting	125 responses
9.6%	None of the above	121 responses
8.0%	Deploying business solutions in different geographies	101 responses
4.2%	Using best-of breed cloud-native tools	53 responses

In this question, we let respondents select as many answers as appropriate to their situation. The top 3 answers were different from last year and, most interestingly, deploying or managing Kubernetes-as-a-Service (KaaS) dropped from top to fifth place. Our experts mostly focused on commenting on this year’s top result, rearchitecting legacy software into microservices:

“This coincides with that infrastructure modernisation objective. Developers are working towards rearchitecting legacy solutions into microservices so that they can be run efficiently on a more modern infrastructure stack.”

—Karthikeyan Shanmugam, Digital Solution Architect, HCL

“A lot of vendors are hoping that Kubernetes will simplify their customer experience with their software. That’s a little bit caught in the general proposition that Kubernetes will simplify everything, which we know isn’t true. People are indeed containerizing their solutions, which is shown in the first answer. And I think the second answer tells you just what it feels like to actually work on Kubernetes: you are working with pipelines of software.”

—Mark Shuttleworth, Founder and CEO, Canonical

“It is great to see that people are more and more adopting open source to build solutions. We still see some monolithic designs working well, but those are inherently more difficult to manage and maintain, especially in distributed environments. Splitting monoliths into smaller pieces and relying more on open source removes some of the dependencies that create maintenance problems. Kubernetes, as the unified slice across all kinds of infrastructure, has helped a lot in removing those dependencies.”

—Ihor Dvoretzkyi, Senior Developer Advocate

“If developers can get away with not having to re-architect anything and can take their current monolith and use it in a container image and deploy it, that is what they are going to do. Re-architecting is about knowing your boundaries and modules, and how to separate them.”

—Michael Hausenblas, Solution Engineering Lead, AWS

## 10. What are the top challenges Kubernetes brings to businesses?

What are your biggest challenges when migrating to/using Kubernetes and containers? Select all that apply.

1240 out of 1279 people answered this question (with multiple choice)

48.0%	Lack of in-house skills/limited manpower	595 responses
37.7%	Company IT structure	468 responses
31.9%	Incompatibility with legacy systems	396 responses
29.3%	Difficulty training users	363 responses
24.8%	Security and compliance concerns not addressed adequately	307 responses
21.6%	Integrating cloud-native applications together	268 responses
16.8%	Poor or limited support from platform providers or partners	208 responses
16.5%	Networking requirements not addressed adequately	205 responses
16.4%	Cost overruns	203 responses
15.6%	Storage/Data requirements not addressed adequately	194 responses
14.9%	Observability / monitoring requirements not addressed	185 responses
13.4%	Inefficient day to day operations	166 responses
11.0%	Cloud platforms don't meet needs/expectations	137 responses
10.7%	Lack of flexibility when it comes to address workload	133 responses
0.6%	Other	7 responses

Adopting Kubernetes is not a walk in the park for most people, with Kelsey Hightower being a famous exception. We found this little soundbite from our experts' interview with Mark Shuttleworth to nicely resume the status quo:

Interviewer: “Do you agree with Kelsey’s opinion that people already have 80% of the necessary skills to be productive with K8s?”

MS: “Well you can walk right? You can walk uphill, right? Why haven’t you climbed Mount Everest?”

Here are some other interesting quotes on the topic:

“I think the ethos of the Kubernetes project—which was the seed crystal for the whole cloud native movement—was to define best practices for how things should work in a greenfield implementation. We didn’t, and for good reasons, spend a lot of time thinking about what that might mean for brownfield customers with existing applications. So if you’re a net new, born in the cloud company, it makes complete sense to just start with Kubernetes, microservices, Go, and the whole ecosystem of cloud native stuff. But if you’re one of the 99% of other users, then you will have stuff that you need to bring with you. You’ve got baggage—and Kubernetes is not so good baggage. We’re getting better, but we built Kubernetes like a walled city. As long as you’re within the city, we’ve got excellent plumbing and paved roads. But if you need to venture outside the city—for example, your city is part of a state—then we don’t really have a lot of roads or paths, or good doorways to the rest of the state. You’re free to leave the city, but then you’re on your own.”

—Tim Hockin, Principal Software Engineer, GCP

“I expect “lack of in house skills and limited manpower” to remain the top challenge in the years to come, as the technology landscape is constantly changing. I see this as a very good challenge to solve. People need to constantly upskill and new job roles are being created as the enterprise needs are evolving with the technology.”

—Ihor Dvoretzkyi, Senior Developer Advocate

“I suspect that difficulty training users, which is another way to look at the skills gap, is only going to increase. This is one of the hard issues that have yet to be addressed in Kubernetes. The people who wanted to learn, the early adopters, have already learned. Now it’s everybody else’s turn to learn, but these people don’t really have an incentive and are going to be harder to teach. There may be a line that some people never cross, there may be some parts of the organisation that just don’t want to go there. Unlike previous generational shifts, like UNIX to Linux, there isn’t really an economic thing to push people. At the end of the day, if you feel perfectly productive, running your app without Kubernetes, you will just keep doing it that way.”

—Mark Shuttleworth, Founder and CEO, Canonical


“When people mention the lack of skill as a blocker, the truth is that they are often already in an environment where they are ready to do the next thing but don’t have the infrastructural or organisational support to do so. It is also a matter of buy versus build: when buying a solution and associated service, an organisation benefits from leveraging external resources and skill set without having to build the capability in-house. When building it in house, the organisation can benefit from implementing its engineering discipline, which could be a useful differentiator.”

—Ken Sipe, Cloud Native Computing Foundation, Edward Jones

## 11. Where do you run Kubernetes?

In which environments do you run Kubernetes clusters? Select all that apply.

1232 out of 1279 people answered this question (with multiple choice)



Respondents were able to choose multiple answers to this question — based on the multiple environments they're running. This year's results are almost identical to last year's data with the lion's share belonging to public clouds, followed by on-prem, with solutions such as VMware, OpenStack and bare metal. This wasn't a surprise for our experts:

"Public clouds, bare metal, OpenStack, have consistently been the top answers for the past couple of years in most surveys. Many businesses use public clouds for their production workloads, because of the ease of use, the flexibility and the availability of resources they provide. On-prem deployments grant users the benefit of full control over the setup and potentially lower costs and latency."

—Ihor Dvoretzkyi, Senior Developer Advocate

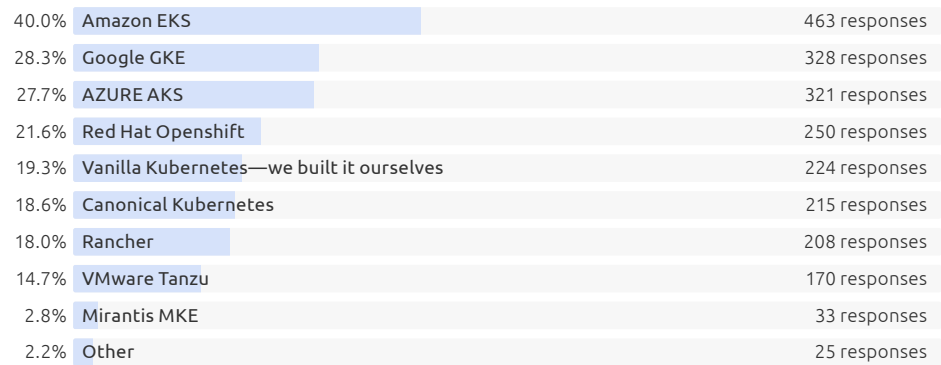
"There's an obvious predominance of public clouds there, but there's definitely more to this picture. It would be interesting to have a closer look at the on-prem deployments, for instance, what some of the bare metal clusters look like."

—Jose Miguel Parrella, Principal Program Manager, Office of the Azure CTO

## 12. K8s distros

Which Kubernetes distribution are you using? Select all that apply.

1158 out of 1279 people answered this question (with multiple choice)



In this multiple answer question, the cloud providers own by far the majority of the Kubernetes distributions market share, which is a natural consequence of their dominance in the K8s cloud environments, as seen in the previous question: users feel more comfortable getting Kubernetes from the cloud they already use. Those who want more than a transactional relationship and a workload scheduler, with custom requirements around hardware and implementation, are looking to run their own clusters.

This trend can be seen with a shift towards distributions like [Canonical Kubernetes](#) that unlock that customisation. Our experts preferred to stay neutral, as they should have, and did not provide any comments on this topic.

## 13. Kubernetes versions and upgrades: too fast or too slow?

What is the oldest version of Kubernetes that you have running on a production cluster?

1184 out of 1279 people answered this question



Last time, our experts expressed their concerns that people were not catching up to the latest fixes and innovations from the K8s community. This year's numbers are a bit better, with a 7% decrease in people that use really really old K8s versions. Is this a sign of maturity in organisations' adoption of the platform? Or a validation of the good work of the community and the vendors? Here is how our experts see this trend:



“The frequency of upgrades depends on the size of your clusters. You don’t want to upgrade giant clusters and may be fine with running a couple of versions behind if they work well for you. At the other end of the spectrum is a huge fleet of smaller clusters. The provider needs to be able to give you an overview of the entire fleet and enable you to upgrade those clusters with short notice.”

—Michael Hausenblas, Solution Engineering Lead, AWS


“I am not equipped to say why people are tolerating to be so many releases behind. It is possible that things have stabilized to a point, so they are ok to stay with a specific version for very long. Overall, I think people are catching up though. Manage clusters have a lot to do with that as well, helping organisations move to the latest releases.”

—Jose Miguel Parrella, Principal Program Manager, Office of the Azure CTO

## 14. Kubernetes in local development

What Kubernetes environment(s) do you target during local container development? Select all that apply.

1222 out of 1279 people answered this question (with multiple choice)



Kubernetes in local development sparked an interesting debate in last year’s report, with one side considering it an unnecessary overhead and the other a nice entry point to a full K8s-based developer workflow. This year’s results confirm Docker’s success in creating a ubiquitous offering and the ease of installation of Kubernetes on top of it, as well as the popularity of Minikube which was the only solution available for users to experiment with K8s for a long time.

On the other hand, some of Docker’s and Minikube’s reported challenges have led to a rise of alternative solutions, such as [MicroK8s](#) and k3s, that keep the simple developer experience but look to a more enterprise-grade, open source and community-led approach.

# Configuration Management

## 15. How many Helm charts has your team used/modified in the last 90 days

How many Helm charts has your team used/modified in the last 90 days?

1194 out of 1279 people answered this question



The number of Helm charts used shows quite similar trends as the previous year, suggesting that not many people are using/modifying them. One reason for this could be that despite helping people to get started with K8s, managing and tweaking Helm charts becomes really complicated, especially at a larger scale. However, there are still a few people that are still using a high number of Helm charts, most likely for custom purposes.

## 16. What is the best way to manage software on Kubernetes?

What is your preferred method for operating, upgrading, and maintaining software on Kubernetes?

1158 out of 1279 people answered this question



Although experts don't think of Helm as an efficient solution for large scale implementations, it is still used by many to manage their software on Kubernetes due to its sheer simplicity to pick someone else's blueprint and get started. Scripts, configuration management tools and GitOps are popular alternatives.

"GitOps is a rapidly emerging trend these days and my prediction is that next year we will see a lot more respondents, between 15% and 20%, selecting that as the answer for this question."

—Ihor Dvoretzkyi, Senior Developer Advocate

“Helm has taken the entire Kubernetes schema, changed it and abandoned the documentation to ask the users to learn a different schema without taking opinions or reducing complexity. I think that in many cases, complexity can be reduced through opinions.

By not holding opinions, you just pass the complexity on to people who are less equipped to make opinions or to have consistent answers.

Actually, DevOps people, who are literally hands-on day-to-day with the applications, don't always know the best practices in terms of norms or monitoring. This is why I'm a big believer in the PAAS pattern – using Kubernetes as a building block and building your own opinionated layers on the top to reduce complexity and improve consistency. Instead of users configuring parameters wrongly, tools like Custom Resource Definitions (CRDs) can help set up specific parameters to homogenise deployments. Even if these things are really just thin veneers over standard Kubernetes abstractions, you do get to make your own opinions on how to do things in ways that Kubernetes can't.


The thing is, Kubernetes tries to fit into a spectrum between Infrastructure-as-a-Service at one end, and Platform-as-a-Service at the other end. Usually, IAAS has no opinions and is oftentimes the Wild West. On the other hand, PAAS can be overly restrictive, and often people can't live with the decisions that it made and/or the opinions that it didn't have. Kubernetes CRDs give users the means to adopt a PAAS model with the possibility to make their own decisions and opinions. I think that's a powerful pattern.”

—Tim Hockin, Principal Software Engineer, GCP

## 17. What is the solution to too many Helm charts?

What is the solution to too many helm charts?

1200 out of 1279 people answered this question




As shown by the 2021 and now 2022 editions of this report, a majority of people are not aware of the solutions to properly handle Helm charts, probably because most organisations are still in the experimental phase of adoption. Operators seem to provide a solution to the issue, as they go deeper into the full lifecycle management and active management, followed by GitOps and other specialised tools.

## 18. In a typical organization, who manages the cloud infrastructure?

### Who manages your organisation's cloud infrastructure?

1241 out of 1279 people answered this question



This question showcases the wide diversity of scenarios when it comes to infrastructure management. While some companies adopt a “you build in, you run it” philosophy, others are implementing a central governance model. When looking at the takes from our experts, it is clear that the responsibility of managing cloud infrastructure seems to be given to various people depending on the type of cloud infrastructure and the expertise needed to manage it. Although many organisations seem to run their own Kubernetes distribution, the adoption of managed services has increased compared to the previous year.

Relevant quotes:

“For organisations of a certain size, the very clear pattern is having a dedicated platform team or a central IT function to manage infrastructure. This platform team might be running their own Kubernetes distribution or consuming the managed Kubernetes service from a cloud provider. This provides a nice abstraction layer for developer teams to interact with. Developers are able to work with a dedicated cluster or namespace provided by the platform team without knowing what is happening under the hood, where the Kubernetes cluster lives, where the storage comes from, etc. That is the true benefit of having a platform team.”

—Michael Hausenblas, *Solution Engineering Lead, AWS*

“I love that “our intern” has lost some ground compared to last year! I’m also pretty sure there are organisations that have self-served Kubernetes clusters for development.”

—Jose Miguel Parrella, *Principal Program Manager, Office of the Azure CTO*

## 19. Top considerations for DevOps in cloud native technologies

What are the TWO most important questions that ops people should care about?

1265 out of 1279 people answered this question (with multiple choice)

37.8%	How secure is that thing?	478 responses
25.5%	How can we optimize resource utilization?	322 responses
20.9%	Is this thing working as it should?	264 responses
19.4%	Which software components should be included in the workflow	245 responses
19.0%	What resources should be allocated to the scenario?	240 responses
17.4%	What's happening in that YAML file?	220 responses
15.4%	Where should the application run?	195 responses
14.7%	How can we optimize applications to reduce work/errors in Day 2 operations?	186 responses
13.7%	What's the fastest way to deploy X?	173 responses
11.2%	What is happening in that helm chart?	142 responses
5.1%	How should that scenario be integrated into the wider estate?	65 responses

While security and resource utilisation remain the top two considerations, optimising applications to reduce efforts for Day 2 operations has dropped significantly on our polls. This might be because we added a few options that focus on more basic challenges, such as “is this thing working as it should”. As per Alexis Richardson’s insight from the previous report “When people worry about the basics [...] this is an indication that they haven’t yet crossed the chasm of full enterprise adoption.”, something that naturally wouldn’t have changed in the span of 6 months between the two surveys.

**“Security remains the key challenge. It is possible that resource utilisation ranked higher this time because you had more infrastructure architects surveyed.”**

—Karthikeyan Shanmugam, Digital Solution Architect, HCL

“Security is not only about code, it’s much deeper than that. And when it comes to operations, security flaws can be extremely damaging. The scope and the importance would explain why it’s common to see “security” as the first thing that comes into people’s minds for these sorts of questions.

For Kubernetes and containers, the biggest questions are around scaling, eviction, and upgrade policies. “How are you going to keep Kubernetes updated?”. Those are, mostly for people running these K8s in production, a day-30 discussion. People are starting to realise there’s no “let’s deploy and see how it goes”. There are several reliable releases a year that include API changes — people know they have got to stay on top of that.”

—Jose Miguel Parrella, Principal Program Manager, Office of the Azure CTO

# CNCF Landscape

## 20. Upskilling in cloud native tech

How do you prefer upskilling in cloud native technologies

1239 out of 1279 people answered this question

27.3%	Following online training courses (MOOCs)	338 responses
20.6%	Google is my friend	255 responses
15.0%	Experimenting on a local K8s cluster	186 responses
10.6%	Going through official certification material	131 responses
10.3%	Watching YouTube videos	127 responses
7.8%	Reading tech blogs	97 responses
4.3%	Shadowing my senior developer	53 responses
3.8%	Reading e-books / listening to audiobooks	47 responses
0.4%	Other	5 responses

While online training courses are the way to go to upskill in cloud native technologies, people also heavily rely on Google to learn more about the topic. The overall trend just confirms the need for training and the growing adoption of cloud native technologies.

**"We've established from the last report that the skills gap is the no.1 problem. The percentage of people doing online training and watching youtube videos to upskill is surprisingly large."**

—Alexis Richardson, Founder and CEO, Weaveworks

**"I think there is a big trend there, with Kubernetes moving more and more towards production, where many people need to upskill in cloud native technologies. Online training courses are great for those who are just starting and want to get a significant amount of knowledge in a determined time frame. For people that are already using Kubernetes and want to answer specific questions or to learn a specific topic in detail, googling information or experimenting hands-on is the preferred approach. I would like to see more people consuming books to get upskilled, this is my preferred way to learn."**

—Ihor Dvoretzkyi, Senior Developer Advocate

## 21. Cloud native technology mastery

Which cloud-native solution space(s) have you mastered? Select all that apply.

1243 out of 1279 people answered this question (with multiple choice)

30.9%	Continuous Integration & Delivery	384 responses
25.6%	Database	318 responses
21.3%	None — I'm pretty good at some of them, but master of none	265 responses
21.2%	Application Definition & Image Build	264 responses
20.8%	Automations & Configuration	258 responses
20.9%	None — I'm still learning	249 responses
18.6%	API Gateway	231 responses
16.6%	Container Registry	206 responses
15.8%	Container Runtime	197 responses
13.6%	Security & Compliance	169 responses
13.5%	Cloud Native Storage	168 responses
13.2%	Observability	164 responses
13.0%	Cloud Native Network	161 responses
12.2%	Streaming & Messaging	152 responses
11.5%	Ingress/Service Proxy	143 responses
11.3%	Scheduling & Orchestration	140 responses
10.9%	Serverless/Function-as-a-Service	135 responses
9.5%	Chaos Engineering	118 responses
8.0%	Service Mesh	99 responses

Our intention with this multiple-choice question was to assess people's comfort in using the different tools from the various categories of the CNCF landscape. The numbers haven't changed much from last time: CI/CD still leads, and Databases are now in second place. It is worth noting that despite the fact that the data indicate that people are mostly still learning, the percentage of people reporting no expertise whatsoever with cloud native tools is evidently decreasing. Our experts agree: mastering cloud native technologies is more complex than originally expected, as they keep evolving and steadily changing users' mindset and skillset.

"I think most people are not even sure what they need to know. It's quite a scary world! Like the fact that security and compliance are at the bottom of this skills list. The one thing that people need more and are really nervous about, i.e. security and compliance, very few have a good understanding of."

—Alexis Richardson, Founder and CEO, Weaveworks

"I think the Kubernetes learning curve is much bigger and steeper than people initially understood."

—Mark Shuttleworth

“CI/CD is certainly an essential part of what Kubernetes can do. Interest in other solution spaces may vary based on one’s profession. For example, a security engineer will certainly be drawn to Kubernetes and container security topics. There are other solution spaces, that although might be easier to understand, are mostly abstracted from the end user, so there’s not too much explicit interest in them.”

—Ihor Dvoretzkyi, Senior Developer Advocate

“The move to the cloud requires strong engineering discipline and it requires a different developer mindset. Programmers now need to have broader skillsets, from infrastructure to YAML to networking.”

—Ken Sipe, Cloud Native Computing Foundation, Edward Jones

## Kubernetes operator

### 22. Would you trust an operator built by an expert?

I would be willing to use a MySQL operator developed by someone who has spent the last 5 years maintaining MySQL in production.

1228 out of 1279 people answered this question

57.3%	Yes	704 responses
21.7%	No although I'd run MySQL on K8s, I wouldn't use the operator	266 responses
21.0%	No, I would never run MySQL on K8s	258 responses

The in-house vs off-the-shelf question is ubiquitous throughout the software industry, and operators are no exception. It seems that organisations are willing to rely on external expertise on that matter, which is not too surprising considering that the skills gap is a major issue in the industry. However, our experts rightfully pointed out the necessity to keep security and compliance in mind when looking for software.

“I certainly would want to use something that is verified, trusted, and respected. If there was an official MySQL K8s operator, that would probably be the first I try. I don’t think we are at a place where you can just take things off the shelf and use them.

Nowadays, if you look at enterprise-grade software, you’ll possibly install the software via the operator the vendor provides. I trust that this will do its job well, but I don’t think it provides a complete solution to deploying and managing the software and K8s.”

—Alexis Richardson, Founder and CEO, Weaveworks



“I think there’s an increasing desire to unburden ourselves from retaining in-house experts in the systems that we depend on. And I think it’s dangerous. If you depend on Postgres and you don’t have a Postgres expert or three on staff – and something goes wrong – your time to recovery is unbounded.”

—Tim Hockin, Principal Software Engineer, GCP

“People that have been consistently solving problems, have gained expertise, and even built communities around a piece of software would definitely be the ones I would trust in operating the software.”

—Ihor Dvoretzkyi, Senior Developer Advocate

Our experts also align with the principle that professionals that operate software day in and day out are the ones to trust to do so consistently. Kubernetes operators are an attempt to codify that operational knowledge to make this approach more scalable. This is a very interesting topic, and we used the subsequent questions to get users’ feedback on it.

## 23. Experience with operator

What is your experience with Kubernetes operator?

1228 out of 1279 people answered this question

28.0%	Trying them out is on my to-do list	344 responses
20.0%	Getting comfortable with them	245 responses
16.0%	Using them for apps in production	197 responses
8.8%	Tried once, failed miserably	108 responses
7.8%	We have security concerns	96 responses
7.7%	Other	95 responses
5.9%	Don't need them	72 responses
5.8%	Don't care for them	71 responses

What is your experience with Charmed operator?

1218 out of 1279 people answered this question

28.7%	Trying them out is on my to-do list	350 responses
14.0%	How are these different from the other operators?	171 responses
11.8%	Other	144 responses
8.9%	Getting comfortable with them	108 responses
8.2%	Using them for apps in production	100 responses
8.0%	We have security concerns	97 responses
7.7%	Don't care for them	94 responses
6.9%	Don't need them	84 responses
5.7%	Tried one, failed miserably	70 responses

## What technology would you most like to see a Charmed operator written for?

1207 out of 1279 people answered this question

37.2%	I don't know	449 responses
17.4%	Cloud-native storage	210 responses
15.9%	Database	192 responses
15.7%	Service mesh	189 responses
13.7%	Ingress/Service Proxy	165 responses
0.2%	Other	2 responses

These questions are painting an interesting picture. Between the two surveys, the amount of people that want to try out operators has stayed more or less the same. However, it seems that a number of respondents did manage to try them out with various levels of success, ranging from confident use to security concerns and failed attempts. Our experts gave their interpretation of the situation.

“I think the problem with operators is the investment of time needed to just try them out and the lack of documentation. There’s often not a clear explanation of how the operator handles corner cases – how do they handle data corruption, migration, and schema updates. In order to figure those things out, you need to test them. People just search the web for “Postgres operator”. “Which one should I use? They’re all from people that I don’t know”.

Maybe the rise of supported operators, paid operators, or distro-supported operators would be useful. If GKE came out and said, here is the PostgreSQL operator that we endorse and you should use this one on GKE, I bet more people will use it.”

—Tim Hockin, Principal Software Engineer, GCP

“I expect automation of operations to be another big thing in the years to come. Kubernetes operators or Helm charts bring automation to Kubernetes, so I consider them a distinct group of software. There should be a commonplace, like an App Store, where people publish and consume these artefacts. There will be specific ownership of the artefacts, validation as well as different flavours of them. And the “store” will have the right information for people to choose the right flavour, based on documentation, ratings, and the different publishers.”

—Karthikeyan Shanmugam, Digital Solution Architect, HCL

“There are so many things around software operations that are often ignored or require a lot of effort to the extent of having a dedicated team just to operate a particular piece of software. For example, if you have a time series database that your provider of choice doesn’t offer as a service, then your only option is effectively to run it yourself. In this case, often the best choice is to deploy the database via an operator, a custom controller that really understands the application workflow and takes care of everything. In general, I think we are still in the very early days of Kubernetes operators.”

—Michael Hausenblas, Solution Engineering Lead, AWS

“Operators improve the way we perform cloud operations. They deliver some of the other value propositions that people seem to be validating, like feature velocity, agility, and better practices. I think it comes down to people not understanding why they would benefit from using them: maybe the defaults might be enough or they can make the whole thing behave as they want to at the application level.”

—Jose Miguel Parrella, Principal Program Manager, Office of the Azure CTO

## Kubernetes advanced usage

### 24. Selection criteria for container images

Which of the following do you value the most when selecting a base image of a container image? Choose 3.

1262 out of 1279 people answered this question (with multiple choice)

55.1%	Security — passed vulnerability and malware scanning	695 responses
39.1%	Compliance — with your company policies or a standard	493 responses
38.9%	Stability — long term supported versions	491 responses
35.5%	Provenance — getting the image from a trusted publisher	448 responses
34.2%	Size — lightweight images	431 responses
27.1%	Developer experience — frictionless usability of the image	342 responses
26.8%	Ready-to-use — default packages and tools included	338 responses
24.1%	Price — lowest cost for pulling or running the image	304 responses
19.3%	Base layer — preference for an Alpine, UBI, Ubuntu, etc	244 responses

While security, stability and size were the top concerns for our respondents last year, the results for this year now feature compliance in the top 3, with provenance creeping up at the 4th place. These considerations reflect the growing importance of procurement over unbridled experiments as the adoption grows.

“Security and stability are certainly production considerations, so it looks to me that people are thinking about the right things as Kubernetes adoption in production steadily grows.”

—Ihor Dvoretzkyi, Senior Developer Advocate

### 25. Stateful applications in Kubernetes

Do you run stateful applications in containers?

1226 out of 1279 people answered this question

33.7%	Yes, in production	413 responses
24.2%	No, stateless applications only	297 responses
23.2%	Yes, evaluating	285 responses
18.8%	No, but planning to in the next 12 months	231 responses

Compared to the results from our 2021 report, the percentage of people running stateful applications in production or currently evaluating them has slightly decreased. Our experts do not consider this to be an indication that people are now reluctant to do so, rather an indication of people being more cautious of the implications.

“I know a lot of people who run stateful applications. There’s a lot of uncertainty out there, about not running stateful applications, but for many cases, it works just fine. But you need to understand what your system is doing. Which means you understand how it’s going to fail.”

—Tim Hockin, Principal Software Engineer, GCP


“Stateless workloads on Kubernetes, to me, are a no-brainer. For stateful applications, you always have to ask yourself “Is there an offering for this that is better managed, or do I have to run it myself?”. And then you need to think about the implications of your decision.”

—Michael Hausenblas, Solution Engineering Lead, AWS

## 26. Kubernetes security: how to isolate your applications

How do you separate applications in Kubernetes? Select all that apply.

1212 out of 1279 people answered this question (with multiple choice)



As a concept, securing applications in Kubernetes is still in its infancy. Namespacing is the de facto way to do so, as it gives a basic form of isolation for resources. There are efforts and projects, such as VCluster, aiming to extend this capability further as the community strives for a Kubernetes multi-tenant architecture.

“By definition, security has multiple angles to it. One is perimeter access, the other component is quality. All Day 0 are bugs that were not understood or known are essentially quality issues. You wouldn’t have security issues if you were to eliminate all quality issues and as long as you had the perimeter managed correctly.”

—Ken Sipe, Cloud Native Computing Foundation, Edward Jones

“What I see is that security remains a top concern. There should be more initiatives or best practices introduced. They should be part of product documentation or in case studies. There should be more initiatives in the forthcoming years toward the security angle. Kubernetes is an abstraction, from the application developer’s standpoint.


Likewise, security also should be abstracted from the user and automated on their behalf. DevSecOps need to be more prevalent across the entire stack. DevSecOps is an extension of the DevOps pipeline that introduces security gating. This allows developers to define the environment they are running their software in and what needs to be protected given that environment to ensure the software is running securely.”

—Karthikeyan Shanmugam, Digital Solution Architect, HCL

## 27. Kubernetes like a Pro. Highly-available. Air-gapped. Offline


Are you running High Availability Kubernetes clusters?

1229 out of 1279 people answered this question



Are you running Kubernetes in an air-gapped/offline environment?

1224 out of 1279 people answered this question




One-half of the respondents run high availability Kubernetes clusters while the other half do not. Many respondents seem to use Kubernetes for highly secure, data-sensitive applications – one-third of the respondents are running Kubernetes in an air-gapped/offline environment.

## Edge computing

### 28. In a typical organization, who manages the cloud infrastructure?

What is your edge use case? Select all that apply.

1242 out of 1279 people answered this question (with multiple choice)



What are your requirements when it comes to implementing an edge strategy?  
Select all that apply.

786 out of 1279 people answered this question (with multiple choice)

49.6%	Security and compliance	390 responses
46.8%	Low latency	368 responses
46.4%	Scalability	365 responses
30.3%	Observability	238 responses
27.9%	Resource allocation	219 responses
23.7%	Offline mode experience	186 responses
23.0%	Supported platforms (Linux flavors, Windows hosts, etc.)	181 responses
22.5%	Network support	177 responses
18.1%	Provisioning & management of the edge fleet	142 responses
15.1%	GPUs for AI/ML	119 responses
13.9%	Service mesh support	109 responses
10.2%	Ability to create a 1 or 2 node cluster	80 responses
10.2%	Device "phone home"	80 responses
9.9%	Message brokers support	78 responses
9.2%	Specific communications protocol support	72 responses
6.9%	EPA features (SR-IOV, DPDK, Numa, etc)	54 responses
6.5%	Diversified edge gateway and leaf devices	51 responses

As the adoption of edge technologies continues to grow, so does the number of use cases, powered by the growth of lightweight Kubernetes distribution such as MicroK8s. The main concerns remain the same, with security and compliance, low latency and scalability at the top. Here's what our experts say:

"50% of respondents have an edge use case, that's pretty significant. If you think about it, anything over 25% in technology creates a market."

"Security and latency?! This is telling me that people don't really understand the challenges of edge just yet. Deploying is the hardest bit of it. Provisioning and managing should have been next, but both are towards the bottom of the list!"

—Alexis Richardson, Founder and CEO, Weaveworks

"Retail is one of those markets where every store of every variety needs some amount of technology. I can't be any other way anymore. But they also are running on very thin margins and often don't have any IT expertise. So, the idea of running a Kubernetes cluster in the back of a 7-11 shop is terrifying. Who is going to come and fix it when something goes wrong? We have yet to see really great tech products around the challenges of this space. This is an opportunity, and tech-forward businesses will probably have a leg-up over the rest, but it's not easy or obvious how to do it."

—Tim Hockin, Principal Software Engineer, GCP

## Industry expert bios

*Alexis Richardson*  
*Founder and CEO*



Alexis is the Founder and CEO of [Weaveworks](#), and the former chairman of the TOC for CNCF. Previously he was at Pivotal, as head of products for Spring, RabbitMQ, Redis, Apache Tomcat and vFabric. Alexis was responsible for resetting the product direction of Spring and transitioning the vFabric business from VMware. Alexis co-founded RabbitMQ, and was the CEO of the Rabbit company acquired by VMware in 2010. Rumours persist that he co-founded several other companies including Cohesive Networks, after a career as a prop trader in fixed income derivatives, and a misspent youth studying and teaching mathematical logic.

---

*Mark Shuttleworth*  
*Founder and CEO*



**CANONICAL**

Mark is a global leader in open source and venture philanthropy. He is the CEO of Canonical, which delivers open-source infrastructure, applications and related services to the global technology market. He is the founder of Ubuntu and benefactor of the Shuttleworth Foundation which underwrites pioneering work at the intersection of technology and society. Previously, he founded Thawte, a global leader in cryptographic security and identity, and participated in ISS mission TM-34. He studied finance and information systems at the University of Cape Town, and orbital ballistics in Zvyozdny Gorodok.

---

*Tim Hockin*  
*Principal Software Engineer*



Tim is a Principal Software Engineer at Google, where he works on the Kubernetes, Google Kubernetes Engine (GKE), and Anthos. He has been working on Kubernetes since before it was announced, and mostly pays attention to topics like APIs, networking, storage, multi-cluster, nodes, resource isolation, and cluster sharing. Before Kubernetes, he worked on Google's Borg and Omega projects, and before that he enjoyed playing in the kernel and at the boundary between hardware and software in Google's production fleet.

---

*Michael Hausenblas*  
*Solution Engineering Lead*



Michael is a Solution Engineering Lead in the [AWS](#) open source observability service team. He covers Prometheus/OpenMetrics, Grafana, OpenTelemetry, OpenSearch, and Fluent Bit upstream and in managed services. Before Amazon, Michael worked at Red Hat, Mesosphere (now D2iQ), MapR (now part of HPE), and in two research institutions in Ireland and Austria.

---

*Karthikeyan Shanmugam*  
*Digital Solution Architect*



Karthikeyan (Karthik) is an experienced Solutions Architect professional with about 20+ years of experience in design & development of enterprise applications across Banking, Financial Services, Healthcare and Aviation domains. Currently engaged in technical consulting & providing solutions in the Application Transformation space. Karthik regularly publishes his technology point of view. His articles on emerging technologies (includes cloud, Docker, Kubernetes, microservices, cloud native development etc.) can be read on his [blog](#).

---

*Ken Sipe*  
*Senior Enterprise Architect*



Ken Sipe is a Distributed Application Engineer working on the challenges around container orchestration. Ken is a co-chair of the Operator SDK and a committer on the CNCF sandbox projects KUDO and KUTTL where he's working to improve the Kubernetes experience. Ken is an internationally recognized speaker, receiving the Top Speaker and JavaOne Rockstar award in 2009, and continues to speak and lead projects on such topics as distributed application development, micro-service based architectures, web application security, and software engineering best practices.

---

*Ihor Dvoretzkyi*  
*Senior Developer Advocate*



Ihor Dvoretzkyi is a Senior Developer Advocate at Cloud Native Computing Foundation, focused on Kubernetes-related efforts in the open source community. With a DevOps and Technical PM background, Ihor has been responsible for the projects tightly bound to the cloud computing space, containerized workloads and Linux systems.

---



*Jose Miguel Parrella*  
*Principal Program Manager,*  
*Office of the Azure CTO*



Jose Miguel is an open source advocate at Microsoft focusing on Linux and cloud native technologies. He has been with the company since 2010 and as of 2020 he is working in the Office of the Azure CTO.

Before Microsoft, he helped expand an open source consulting firm as CTO, was member of the technical staff at an electric utility and led many large scale projects as a consultant.

---

## Acknowledgements

The Canonical team wants to thank all the community members that contributed their time, experience, and answers to the survey, which made building this report possible, and to the industry experts who enhanced it with their insightful commentary. We are always keen to get your feedback to improve, and we're happy to ask questions to the community on your behalf. This report is published annually. If you are interested in contributing your voice to it, be on the lookout for the next communication about the survey.

If you liked what you read here, feel free to share the link to its page or to this PDF version.