CANONICAL

Cloud pricing report 2022

Contents

Key findings

- Cost is one of the main drivers behind cloud infrastructure choices. Most teams feel that cloud infrastructure takes a significant portion of their budget.
- When comparing cloud pricing across various providers, what matters most is the total cost of ownership (TCO) - total expenses over a given period of time.
- When it comes to choosing cloud platforms, the majority of cloud specialists make decisions based on TCO calculators.
- While hyperscalers continue to be the most economical option for small-scale application deployments, organisations can reduce their spending by building their own cloud infrastructure once their cloud workloads increase.
- A hybrid multi-cloud architecture proves to be the most economical option in medium and large-scale environments, ensuring maximum CAPEX and OPEX efficiency.

Executive Summary

Estimating and comparing costs between various cloud platforms has always been challenging. While public cloud providers maintain the list prices for their services in the form of cost-per-resource metrics, resembling the same in the private cloud space is not trivial. However, even clear cost-per-resource metrics do not give a complete picture of the total cost of ownership (TCO), as businesses are usually not able to calculate their demands for all types of cloud resources on their own. As a result, leading cloud providers have been maintaining TCO calculators to help their customers estimate costs and make data-driven decisions.

This report consists of three sections. First, we present cloud list prices from leading public and private cloud providers as of July 2022. Then, we demonstrate three sample cost scenarios and estimate the TCO of running the same workload in Amazon Web Services (AWS), Azure, Google Cloud Platform (GCP) and Canonical OpenStack, using TCO calculators provided by those vendors. Finally, we present results from Canonical's Cloud Pricing Survey 2022, along with commentary from industry experts. The survey's outcomes show the growing importance of hybrid multi-cloud architecture and confirm its advantages from a TCO stance.

About Canonical

Ubuntu Pro is a premium Ubuntu image provided by Canonical, delivering the most comprehensive open-source security and compliance for production mission-critical workloads running on public clouds.

Canonical OpenStack is an enterprise cloud platform engineered for price-performance that serves as a cost-effective extension to hyperscale infrastructure.

Canonical Kubernetes is an enterprise multi-cloud, CNCF-certified Kubernetes, allowing for fully composable cluster deployments and streamlined container orchestration.

As the publisher and maintainer of Ubuntu, Canonical looks across the entire spectrum of open source to make sure that whoever chooses Ubuntu as a default platform for their organisation is successful on their digital migration journey. Our mission is to bridge the gap between applications and infrastructure, providing a stable, secure and feature-rich Linux distribution for public clouds, data centres and the edge that works across bare metal, virtual machines (VMs) and containers.

As an inherent part of this strategy, Canonical partners with leading public cloud providers to deliver Ubuntu images optimised for hyperscale infrastructure, including commercial images designed for enterprise customers - Ubuntu Pro. At the same time, Canonical delivers comprehensive products for building an owned infrastructure. Those include Canonical OpenStack and Canonical Kubernetes. We help Ubuntu users optimise their infrastructure spending, making sure that they fully benefit from digital migration.

Cloud list prices

Cloud list prices are what organisations usually check first when estimating the TCO of their workloads. List prices enable them to roughly evaluate costs, compare various cloud providers and make initial decisions regarding the cloud architecture being used.

Public cloud platforms

In order to size up the list price of leading public cloud providers, one has to be opinionated about what exactly to compare. This is because there are hundreds of different instance types available in their portfolio and tens of different metrics that are metered and billed according to usage. Since compute resources usually represent the biggest portion of the TCO, this report only analyses cost-per-instance metrics.

The following criteria were applied to compare cost-per-instance metrics across AWS, Azure and GCP:

- **Pricing model:** on-demand instances
- **Region preference:** US East (Northern Virginia)
- **Operating system:** Ubuntu
- **CPU family preference:** AMD EPYC 7000

The cost-per-instance metrics were collected for general-purpose, compute-optimised and memory-optimised instance types, each with 4 virtual central processing units (vCPUs).

List prices from leading public cloud providers are shown in Tab. 1.

Instance type	AWS instance / price	Azure instance / price	GCP instance / price
General-purpose	m5a.xlarge / $0.172000	D4a v4 / $0.192000	n2d-standard-4 / $0.190320
Compute-optimised	c5a.xlarge / $0.154000	F4s v2 / $0.169000	n2d-highcpu-4 / $0.140508
Memory-optimised	r5a.xlarge / $0.226000	E4a v4 / $0.252000	n2d-highmem-4 / $0.256736

Tab. 1 List pricing for public cloud platforms.

Based on this data, AWS provides the lowest list pricing for general-purpose and memory-optimised instances. Although GCP provides the lowest list pricing for compute-optimised instances, it's important to mention that such instances on GCP use 1 GB of random-access memory (RAM) per 1 vCPU, while the same instance types on AWS and Azure use 2 GB of RAM per 1 vCPU. This allows Google to offer compute-optimised instances at a better price. Besides this, there are no significant pricing differences between leading public cloud providers.

Private cloud platforms

Comparing the reference pricing across leading private cloud providers is much more straightforward. When it comes to private cloud implementation, the vendor's service fee usually consists of the licence costs, design and delivery costs and the ongoing support subscription costs. In addition, some vendors offer fully-managed services for their cloud platforms in order to offload organisations' operations teams.

Service fees from leading private cloud providers are shown in Tab. 2.

Service fee	VMware vRealize	Red Hat OpenStack Platform*	Canonical OpenStack
Licence	$7,945 per 1 CPU or 25 VMs	$0	$0
Design and delivery	$587,000 per cloud	~$10,000 per week of engagement	$75,000 or $150,000 per engagement
Support only	$1,986 per 1 CPU or 25 VMs per annum	~$6,300 per 2 CPUs per annum	$1,500 per host per annum
Support and fully-managed	Not available	Not available	$5,475 per host per annum

Tab. 2 Service fees for private cloud platforms.

* - Red Hat OpenStack Platform list pricing is not publicly available. The provided prices are indicative and are based on third-party reports and sources.

Based on this data, prices and pricing models differ substantially across leading private cloud providers. First of all, VMware is the only vendor that requires licences to be purchased for their vRealize platform. While both VMware and Canonical offer consulting services at a fixed price, Red Hat charges based on the duration of the engagement - which is generally between several weeks and months. When it comes to support subscriptions, VMware and Red Hat make them dependent on the number of CPUs or VMs, while Canonical charges based on the number of hosts. Canonical is also the only vendor that offers fully-managed services for their OpenStack platform.

TCO calculators

Although cloud list prices help organisations estimate costs, a lot of other factors have an impact on the total cost of ownership (TCO). This is because additional charges are applied based on storage and network consumption; significant discounts are available when running workloads in the long-term. In the case of private clouds, hardware costs, hosting charges and operations and maintenance costs should be included for comprehensive estimates. Therefore, when comparing prices across various cloud providers it's best to use TCO calculators. The following section provides a brief overview of the most popular ones.

AWS TCO calculator

AWS' TCO calculator provides TCO estimates of cloud instances based on their number, configuration and additional requirements regarding storage and network [1]. Various regions, guest operating systems and pricing models are available to enable more detailed estimates. The calculator estimates total monthly and yearly costs.

Azure TCO calculator

Azure's TCO calculator provides similar capabilities. It helps estimate the TCO of cloud instances based on their number, configuration and additional requirements [2]. What it lacks, compared to the AWS TCO calculator, is an advanced estimate option that specifies the workload and the ability to assess more granular reservation options.

GCP TCO calculator

GCP's TCO calculator is very similar to the Azure TCO calculator in terms of capabilities [3]. Again, it calculates the TCO of cloud instances based on their number, configuration and requirements for storage and networking. Much like the Azure TCO calculator, it misses some more advanced features compared to the AWS TCO calculator.

Canonical OpenStack TCO calculator

Canonical OpenStack TCO calculator estimates the TCO of Canonical OpenStack cloud platform based on the number of instances, their configuration and required support level [4]. More accurate estimates are available on-demand based on detailed requirements. All estimates contain typical CAPEX and OPEX costs, including reference hardware costs, hosting charges and operations and maintenance costs.

Sample cost scenarios

In the following section, we demonstrate sample scenarios and provide guidance on cloud architecture choices. All estimates were obtained using TCO calculators presented in the "TCO calculators" section, augmented with the list pricing presented in the "Cloud list prices" section wherever needed. All public cloud estimates assume a reserved instances pricing model with a 3-year timeframe and Ubuntu guest OS. Canonical OpenStack estimates assume a 3-year hardware renewal period and include reference hardware costs, typical hosting charges, operations and maintenance costs and fully-managed service.

Scenario 1: Small-scale - Internal CRM system

In this scenario, a software company is hosting a third-party customer relationship management (CRM) system for internal purposes. The platform consists of a database and web application. The estimated requirements for this system are shown in Tab. 3.

Purpose	Number of vCPUs	Amount of RAM [TB]	Amount of persistent storage [TB]	Network requirements	Additional information
Database	24	0.2	2	No	No
Web application	24	0.1	0	No	No

Tab. 3. Small-scale sample requirements.

The estimated TCO for this system across leading public cloud providers and Canonical OpenStack is shown in Tab. 4 and Fig. 1.

Cloud platform	CAPEX [USD]	OPEX annual [USD]	TCO [USD]
AWS	0	12,819.36	38,458.08
Azure	0	11,869.80	35,609.40
GCP	0	11,102.16	33,306.48
Canonical OpenStack	420,124	143,700	851,224

Tab. 4. Small-scale TCO estimates.

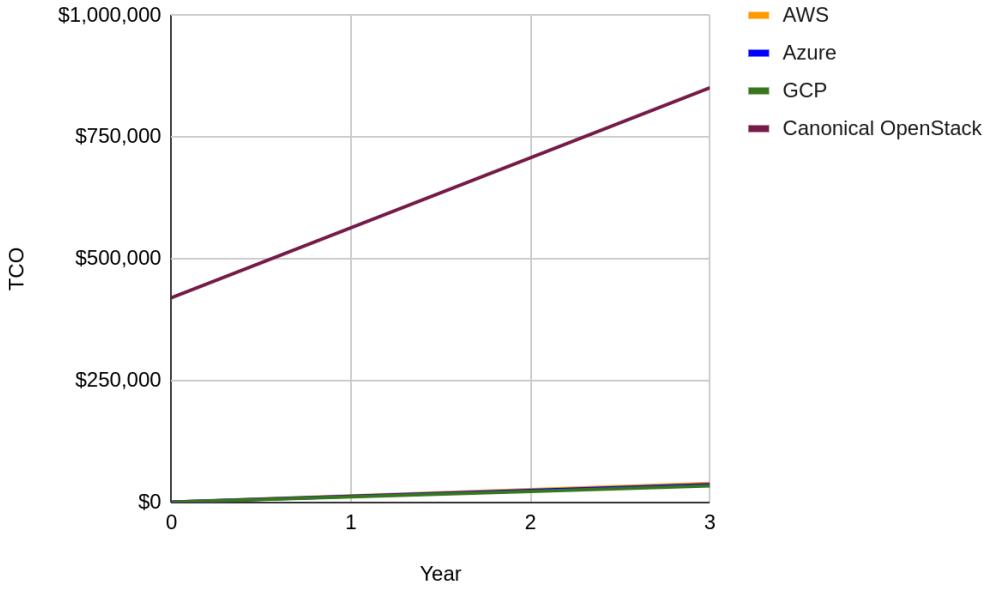

Fig. 1. Small-scale TCO estimates.

In this case, there is no economic justification for private cloud implementation. The organisation should host all workloads in the public cloud, but continue monitoring resource consumption and consider moving to a private cloud if the number of workloads grows.

Scenario 2: Medium-scale - Online banking system

In this scenario, a financial institution is hosting its own online banking system in the cloud. The system consists of a database, web application and a number of developer VMs. The estimated requirements for this system are shown in Tab. 5.

Purpose	Number of vCPUs	Amount of RAM [TB]	Amount of persistent storage [TB]	Network requirements	Additional information
Database	500	4	40	Moderate inter-AZ traffic	No
Web application	1,000	4	0	Moderate inbound traffic	Daily 16-hour load spike
Developer VM	200	0.8	5	No	Used during business hours only
Developer VM	20	0.08	0.5	No	ARM CPUs required; Used during business hours only

Tab. 5. Medium-scale sample requirements.

The estimated TCO for this system across leading public cloud providers and Canonical OpenStack is shown in Tab. 6 and Fig. 2.

Cloud platform	CAPEX [USD]	OPEX annual [USD]	TCO [USD]
AWS	0	472,247.46	1,416,742.38
Azure	0	391,502.04	1,174,506.12
GCP	0	361,619.04	1,084,857.12
Canonical OpenStack	420,124	143,700	851,224

Tab. 6. Medium-scale TCO estimates.

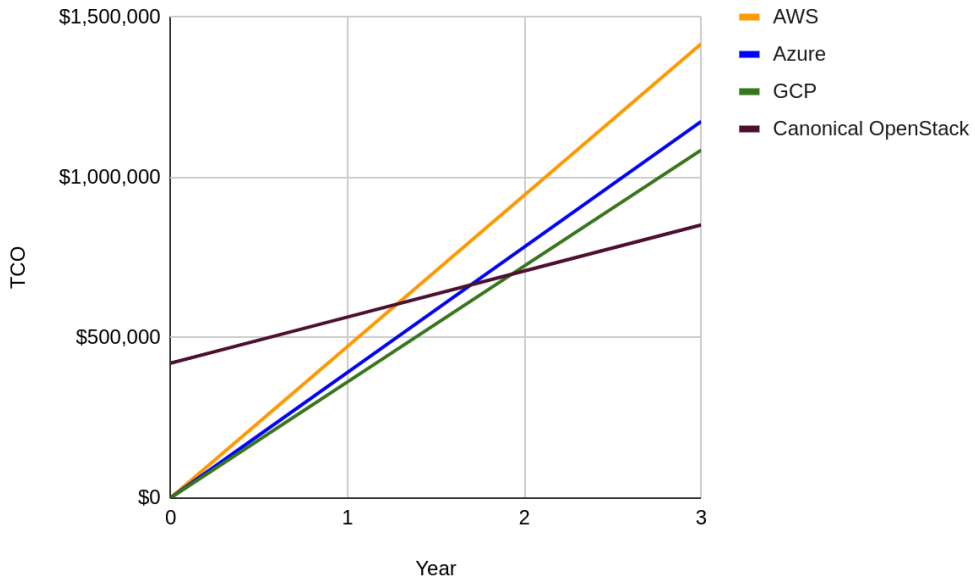

Fig. 2. Medium-scale TCO estimates.

In this case, it makes more sense to deploy cost-effective private cloud infrastructure and host the majority of the workloads there. As the implementation of additional ARM-based hypervisors for the purpose of hosting a small number of developer VMs is not economically feasible, the organisation can use hybrid multi-cloud architecture to host those instances in the public cloud.

Scenario 3: Large-scale - Video streaming system

In this scenario, an entertainment company is hosting their own video streaming system. The system consists of a data warehouse, data lake, data analytics, video transcoding engine and web application. The roughly estimated requirements for this system are shown in Tab. 7.

Purpose	Number of vCPUs	Amount of RAM [TB]	Amount of persistent storage [TB]	Network requirements	Additional information
Data warehouse	32,000	256	5,120	Moderate inter-AZ traffic	No
Data lake	4,000	32	256	No	No
Data analytics	12,000	24	0	No	Run daily for 8 hours
Video transcoding engine	24,000	48	0	No	Daily 8-hour load spike
Web application	6,000	24	0	High outbound traffic	Daily 8-hour load spike

Tab. 7. Large-scale sample requirements.

The estimated TCO for this system across leading public cloud providers and Canonical OpenStack is shown in Tab. 8 and Fig. 3.

Cloud platform	CAPEX [USD]	OPEX annual [USD]	TCO [USD]
AWS	0	26,020,604.88	78,061,814.64
Azure	0	24,167,781.72	72,503345.16
GCP	0	22,038,250.08	66,114,750.24
Canonical OpenStack	8,802,259	3,249,625	18,551,134

Tab. 8. Large-scale TCO estimates.

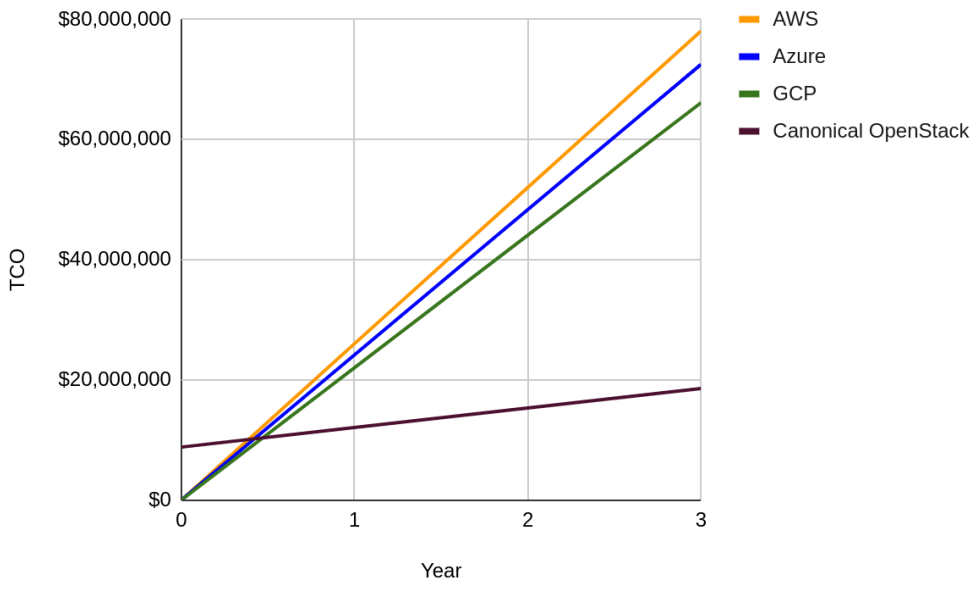

Fig. 3. Large-scale TCO estimates.

In this case, a hybrid multi-cloud architecture is a must as there is no economic justification for running workloads in the public cloud. The organisation can leverage hyperscale infrastructure during heavy load periods while hosting the majority of its workloads in their private cloud. At this scale, it also makes more sense to hire a dedicated cloud operations team rather than relying on fully-managed services.

Hybrid multi-cloud architecture

The scenarios presented above indicate that while public clouds tend to be more economic for small-scale deployments, investing in private cloud infrastructure pays off when running workloads at a larger scale. However, since the nature of many business applications is rather dynamic, hybrid multi-cloud architecture proves to be the most optimal approach in the long-term. As a result, it should be an integral part of the digital migration strategy.

Implementing hybrid cloud architecture enables organisations to optimise their infrastructure costs by using both private and public clouds at the same time, gaining the benefits of both. Businesses can run the majority of their workloads in the private cloud and only use highly scalable, but more expensive public cloud resources when absolutely necessary. This includes handling tasks during heavy load periods as well as running occasional resource-intensive workloads, such as data analytics.

In turn, implementing multi-cloud architecture - using at least two cloud platforms from two different cloud providers at the same time - allows them to avoid vendor lock-in as they are no longer dependent on a single provider. It also enables them to improve the reliability of their workloads, as distributing business applications across multiple cloud platforms eliminates the negative impact of potential cloud outages.

Canonical's Cloud Pricing Survey 2022

In order to better understand cloud pricing concepts and their impact on cloud infrastructure choices, Canonical ran a community survey to collect responses from hundreds of cloud users. Canonical's Cloud Pricing Survey 2022 was hosted from the 7th of Apr 2022 to the 31st of May 2022 and consisted of 40 questions. 326 people participated in the survey. In the following chapter, we present survey results along with a commentary from industry experts.

1. I am best described as (choose one):

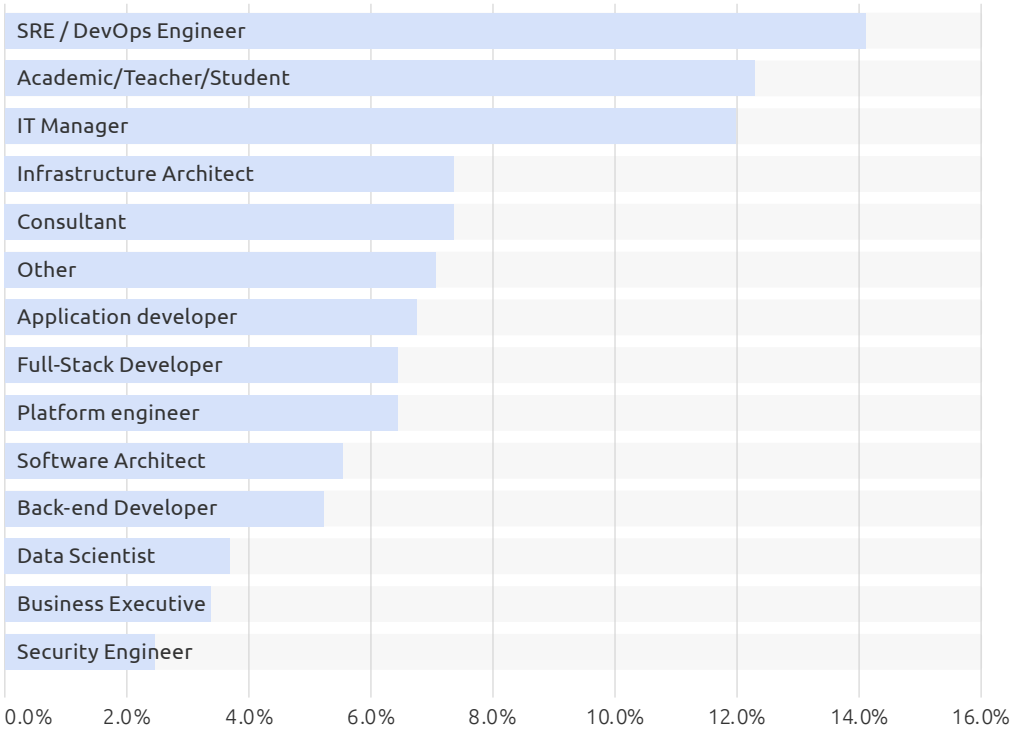

"The shift to become a digital business can clearly be seen in the roles people in IT are now doing in greater numbers. DevOps/SRE is the top response, with 14% of respondents classing themselves in this way".

– Roy Illsey, Chief Analyst, Omdia

2. My company's industry is best described as (choose one):

"It's interesting to see that most people in our survey work in software and technology. It could be indicative that cloud cost optimisation is more of a priority in software and technology compared to other industries which may be lagging behind in digital transformation".

– Tytus Kurek, Product Manager, Canonical

3. What are you using the cloud(s) for? (choose all that apply):

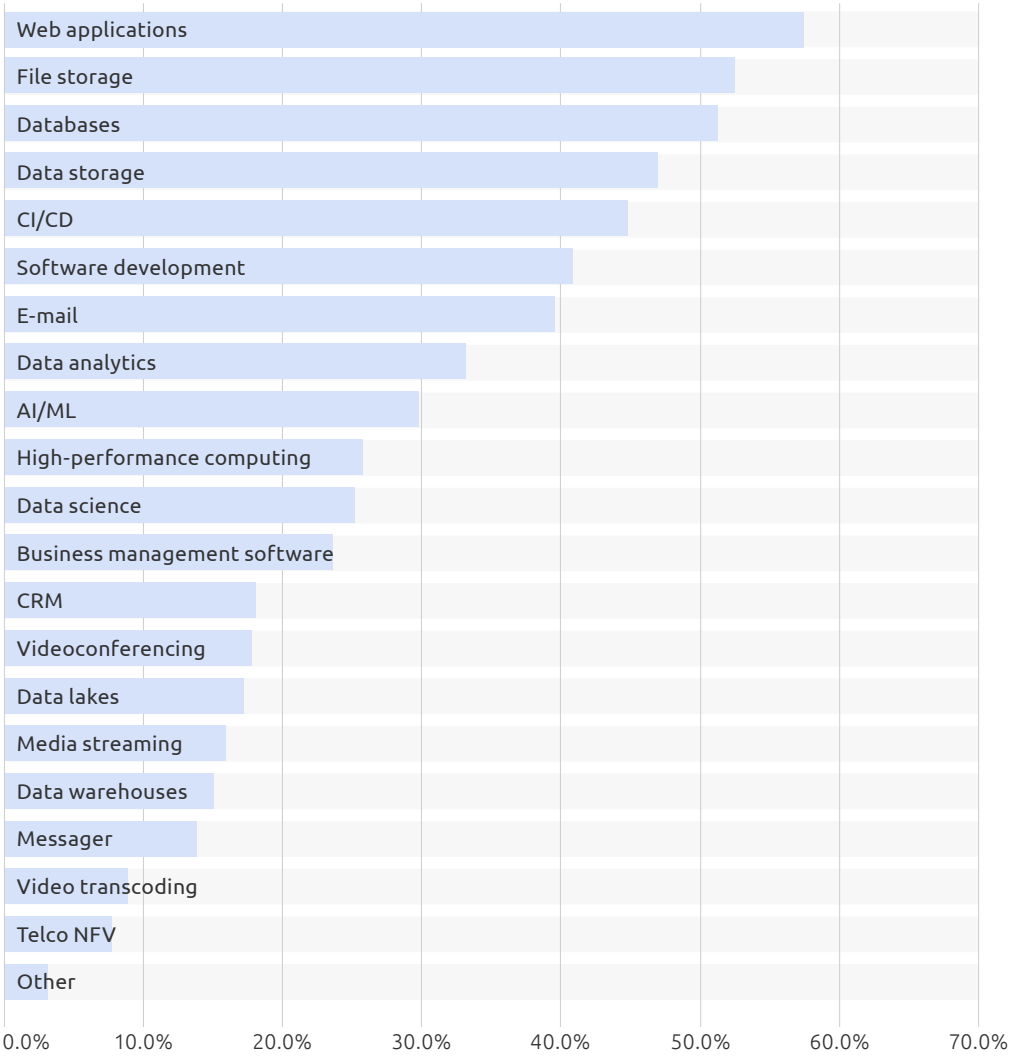

"The wide range of use cases shows that innovation is still happening in the cloud infrastructure space, and requirements are always evolving. An openly-developed open source solution like OpenStack allows us to capture those evolving requirements and new use cases. OpenStack is about to deliver its 26th on-time release, and the project continues to attract upstream engagement from thousands of developers. Their consistent work over the arc of 12 years has evolved the project to embrace new features that accommodate new use cases and extend the software's flexibility for existing ones".

– Thierry Carrez, General Manager, Open Infrastructure Foundation

"Three of the top 4 use cases for cloud are data related, this is interesting as the rise of data privacy and sovereignty are becoming more of an issue for nations /regions. This use of the cloud will force the CSPs to adopt the concept of sovereign cloud".

– Roy Illsey, Chief Analyst, Omdia

"Frankly, is there a mainstream organisation *not* hosting its videoconferencing in the cloud? On-site CRM and email are increasingly uncommon, too".

– Bryan Betts, Principal Analyst, Freeform Dynamics

4. Which of the following ones best describes the cloud (choose one)?

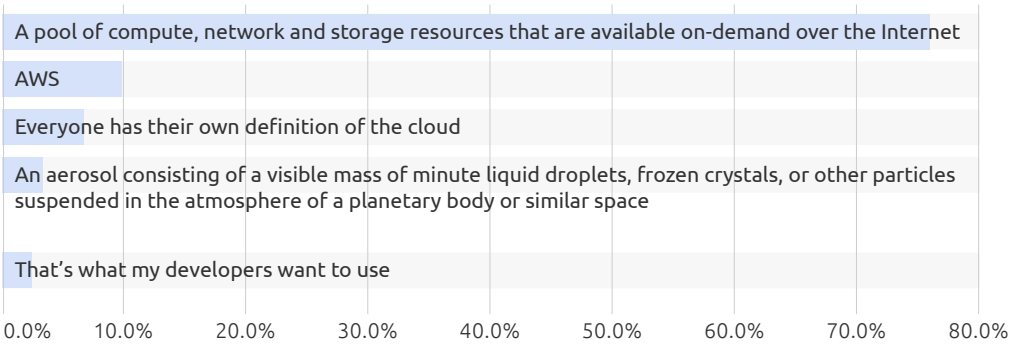

"There's no single definition of the cloud, but the top one really summarises it well. Yet, it was Amazon who set de facto standards for cloud computing so I'm not surprised to see that many think AWS is synonymous with cloud".

– Tytus Kurek, Product Manager, Canonical

5. How long has your organization been using the cloud (choose one)?

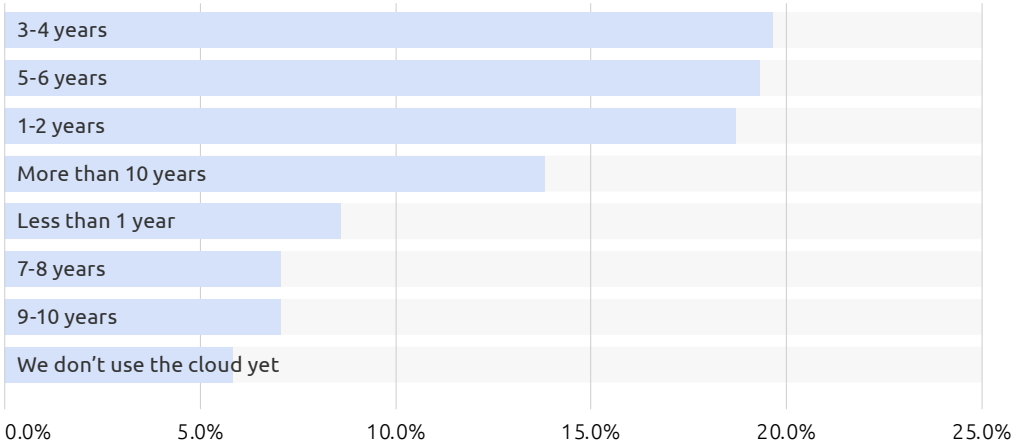

"While cloud has been around for at least 15 years, numerous research reports—and the insights of hyperscale cloud provider executive leadership—show that the largest portion of cloud growth remains in front of us, and there remains a significant market opportunity to migrate workloads to the cloud".

– Thierry Carrez, General Manager, Open Infrastructure Foundation

6. I am consuming cloud services mostly in the form of (choose one):

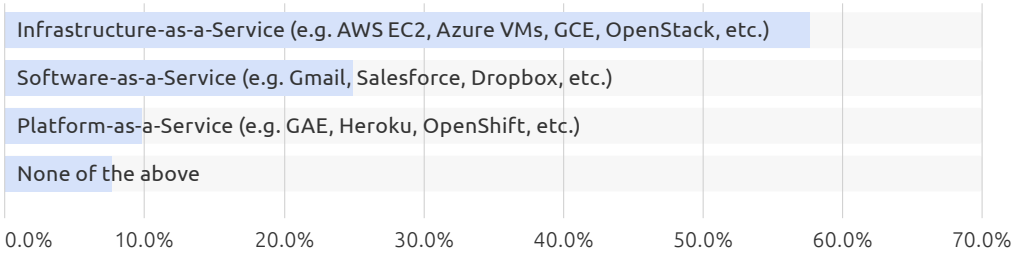

"Even though the majority of respondents consume cloud services as IaaS, the shift towards SaaS is evident. It is going to take years until it becomes the default model, but clearly more and more organisations are starting to explore it. The last few years were spent on multi-cloud integration. Multi-SaaS integration is most likely going to be another challenge".

– Tytus Kurek, Product Manager, Canonical

7. Where do you run your business applications (choose one)?

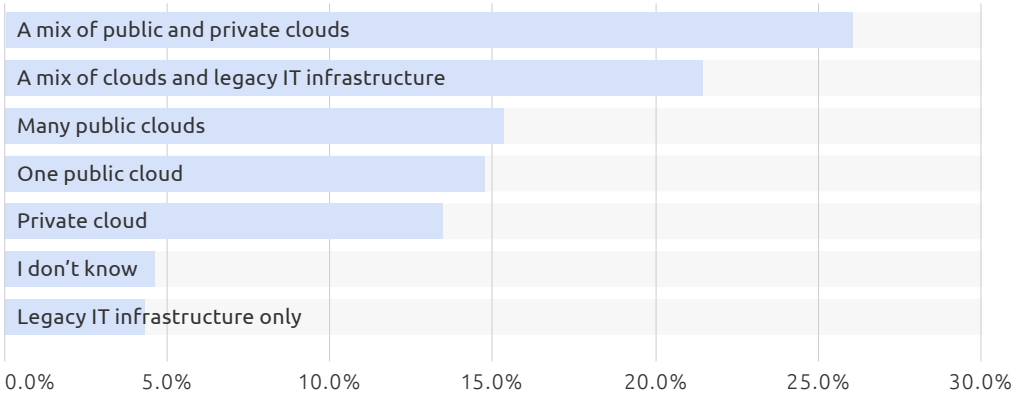

"A mix of clouds is undoubtedly the reality of today. The cloud computing landscape is broad. It includes hyperscalers, hundreds of local public cloud providers and many private cloud platforms, including proprietary and open-source ones. It is not a surprise then that the majority of the respondents use more than one of them at the same time.

Another interesting fact is that more than 25% of those organisations haven't completed their cloud migration yet. It stresses the fact that cloud migration is a long process that's going to continue for the next few years".

– Tytus Kurek, Product Manager, Canonical

8. Which public cloud platforms do you use (choose all that apply)?

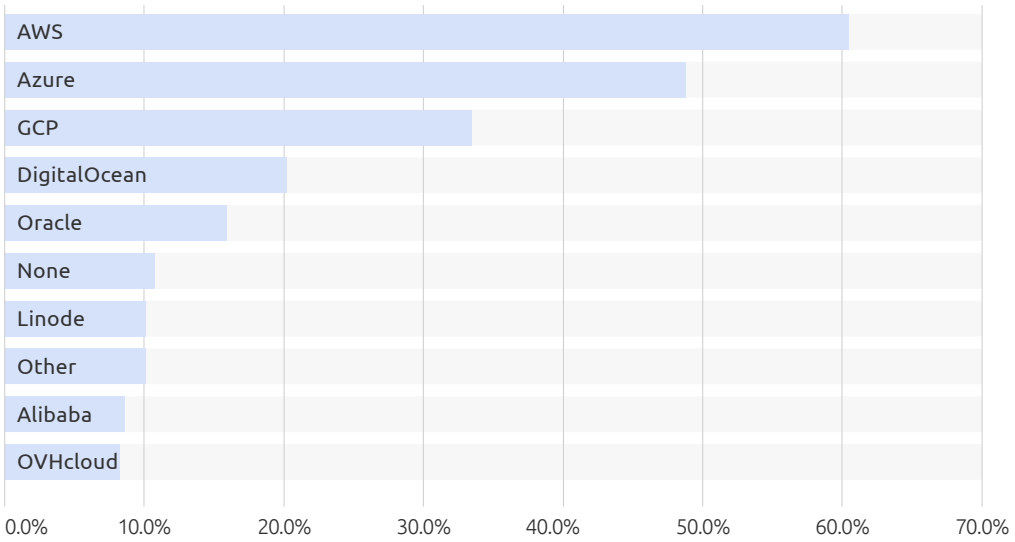

"This result agrees with Omdia's research in terms of the top 3 public cloud providers being used. The interesting insight is that the data shows on average respondents are using 2.3 of these clouds, which demonstrates that multi-cloud is a reality, even if these multi-clouds are not interoperable."

– Roy Illsey, Chief Analyst, Omdia

9. Which private cloud platform do you use (choose all that apply)?

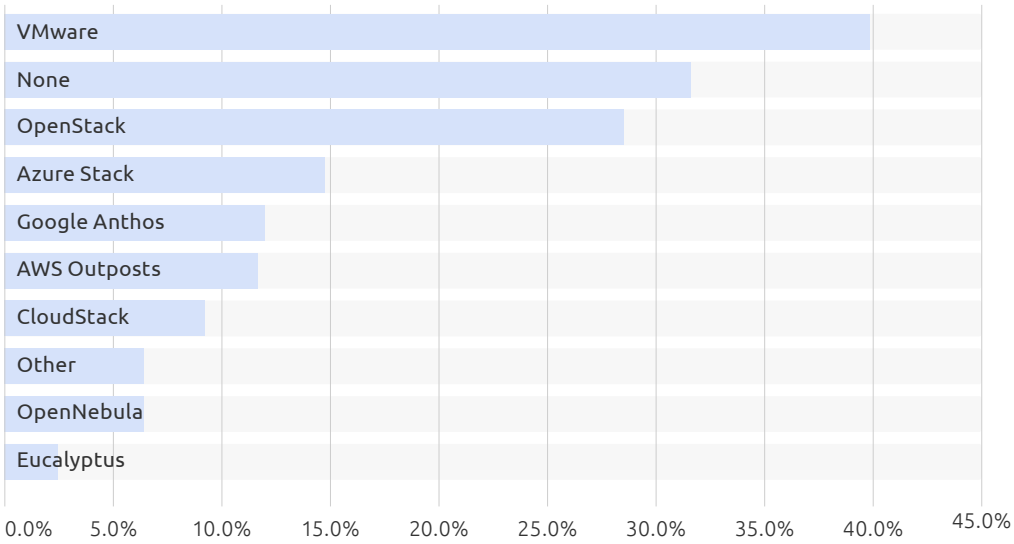

"When organisations look at open source options for private cloud, OpenStack comes out at the top of the list every time. It remains the de facto open source private cloud standard".

– Thierry Carrez, General Manager, Open Infrastructure Foundation

10. How many bare metal machines do you run in total (choose one)?

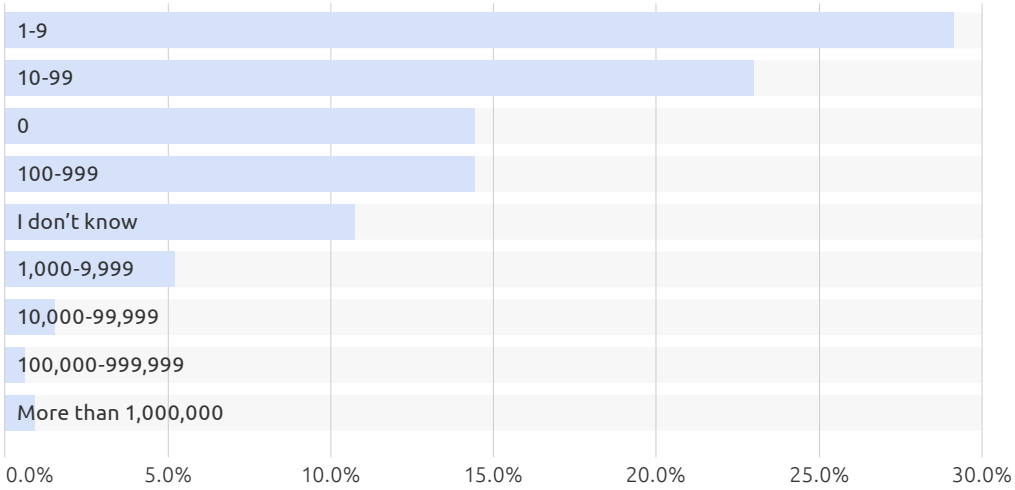

"This data point shows that the majority of those environments are pretty small. However, a cloud with 10-99 bare metal machines can still span across multiple racks. 0 might indicate that all workloads run in the public cloud. It is interesting to see that more than 10% of respondents have no idea how many physical machines they run in total".

– Tytus Kurek, Product Manager, Canonical

11. How many virtual machines do you run in total (choose one)?

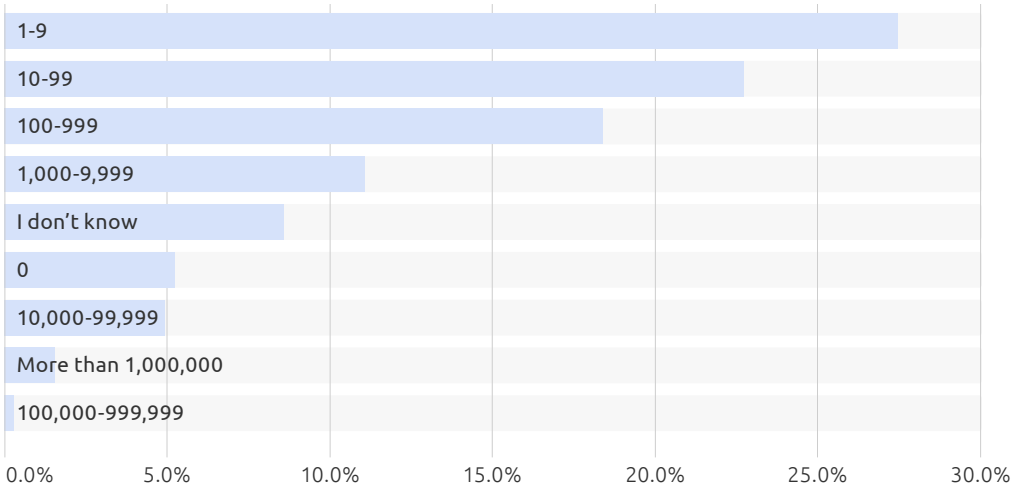

"This graph looks almost like the Pareto distribution. The bigger the environment, the fewer environments like that in total.

0 might mean two things here: either it's all SaaS or it's all containers. It would be interesting to see how many of those organisations have fully moved to cloud-native already".

– Tytus Kurek, Product Manager, Canonical

12. How many containers do you run in total (choose one)?

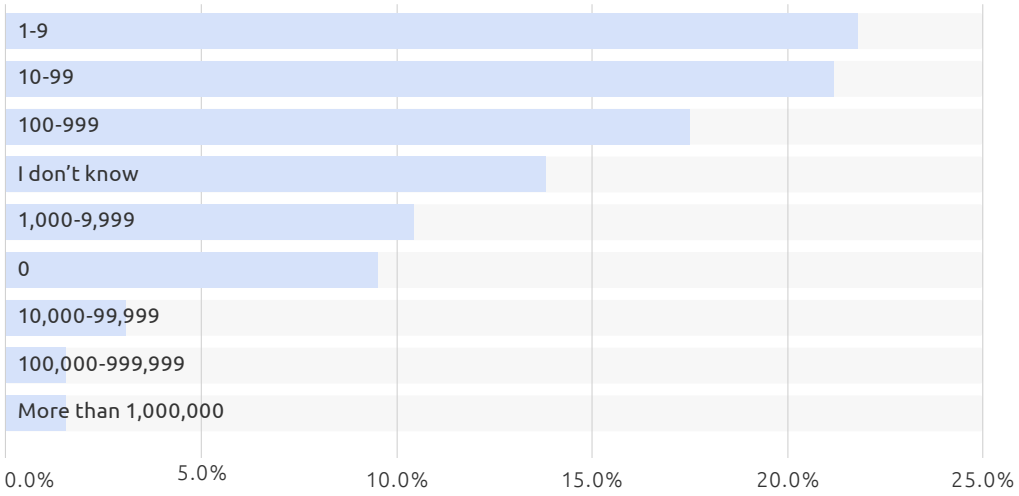

"As the granularity of workloads grows, it becomes more challenging to track all the resources being used. I'm not surprised to see that almost 15% of respondents don't know how many containers they run in total. I often get lost when experimenting with Kubernetes on my workstation. Not to mention big production environments".

– Tytus Kurek, Product Manager, Canonical

13. How much storage do your workloads consume in total (choose one)?

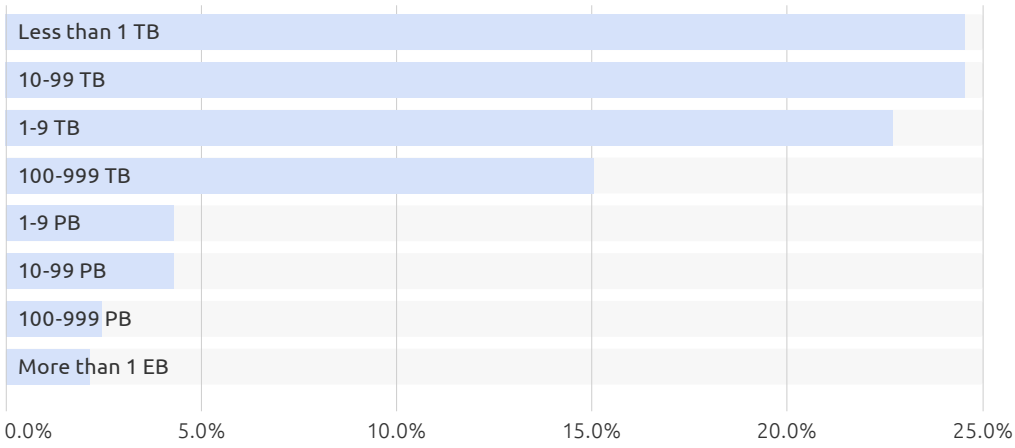

"The amount of storage used for all workloads seems very small, with almost 48% reporting less than 10 TB, Omdia does not believe this is a true usage figure, as when backups and snapshots are taken into account this number could increase significantly".

– Roy Illsey, Chief Analyst, Omdia

14. The craziest setup I have ever tried was (choose one)

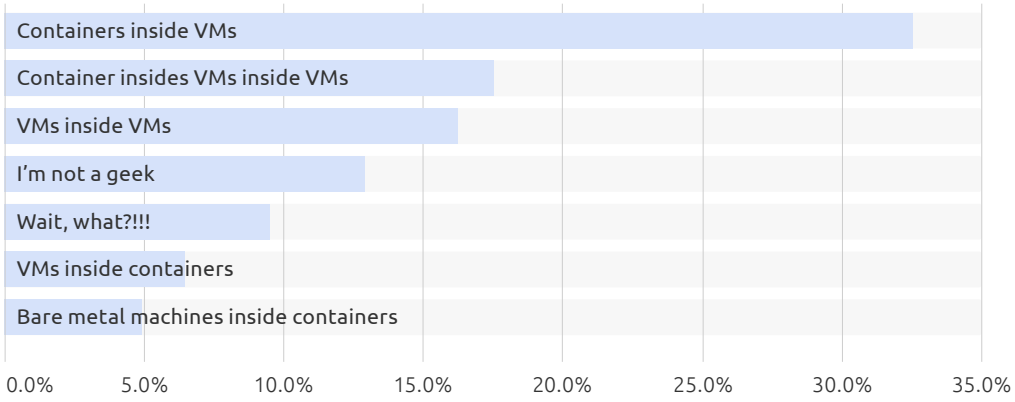

"This question sounds funny, but it really makes me think. The complexity level of modern IT infrastructure can be really high. Detailed tracking of all resources used can be quite a challenge. The same applies to ensuring integration between individual workloads".

– Tytus Kurek, Product Manager, Canonical

15. Who built your cloud infrastructure (choose one)?

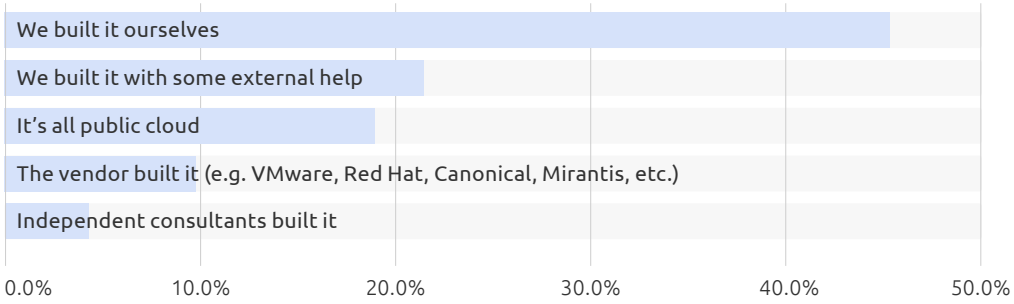

"This question breaks some stereotypes. Cloud technologies are typically considered to be complex. Meanwhile, it turns out that most organisations managed to build their cloud infrastructure on their own. It would be interesting to see the correlation between individual respondents and the specific technologies they use".

– Tytus Kurek, Product Manager, Canonical

16. How do you operate your cloud infrastructure (choose one)?

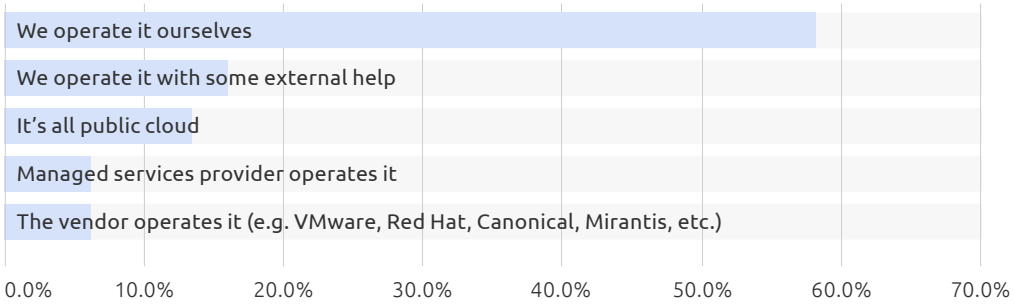

"This is even more reassuring. If people can't use the product, perhaps the product is broken. Here it's evident it's not".

– Tytus Kurek, Product Manager, Canonical

17. Which DevOps tools do you use for cloud infrastructure operations (choose all that apply)?

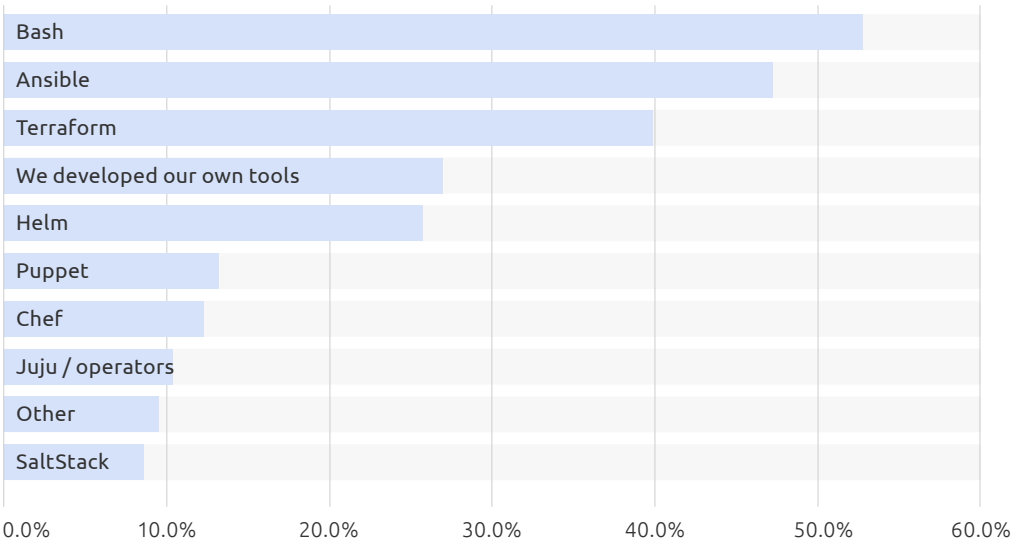

"We've been in a long journey with all these tools. Decades later, Bash still beats both Ansible and Terraform. Is it because of the flexibility it brings? Or is it because this is what everyone learns during one of their first Linux classes? Some legacy tools, like Puppet, Chef and Saltstack are letting us down these days. In the Kubernetes space, Helm continues to lead. However, the rise of operators is noticeable".

– Tytus Kurek, Product Manager, Canonical

18. What is the size of your cloud infrastructure operations team?

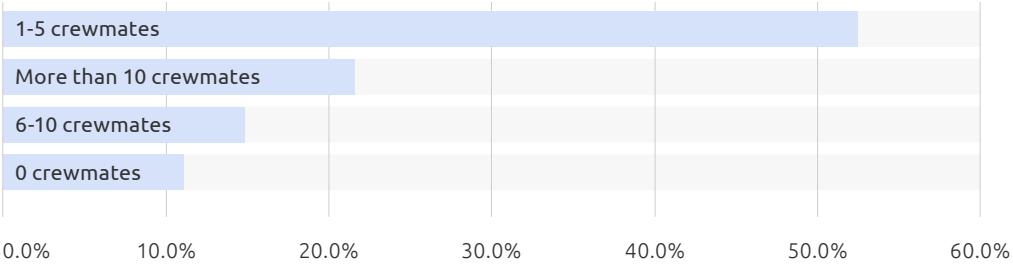

"This is a tricky challenge. We know that operating the cloud 24/7 with a very high SLE requires around 10 FTEs. Does it mean that the majority of cloud environments don't need such a high service level?

The last answer is well-correlated with question 16 where some said that their cloud environment is operated by a managed service provider or a vendor".

– Tytus Kurek, Product Manager, Canonical

19. Has any of those ever happened to you (choose all that apply)?

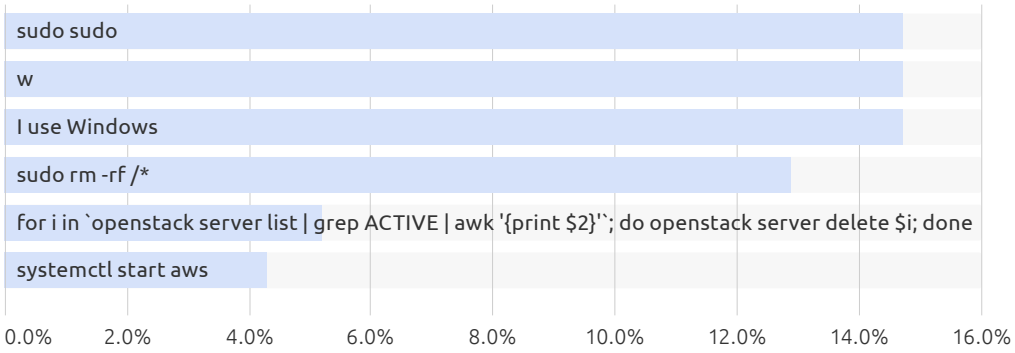

"As we've seen in question 17, Bash continues to be the leading DevOps tool. Not surprisingly, a lot of respondents have used some of those commands then. My compassion goes out to those who accidentally removed rootfs".

– Tytus Kurek, Product Manager, Canonical

20. Which is most important to you (choose one)?

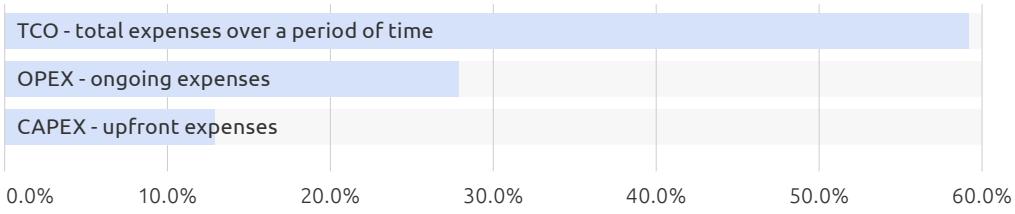

"The metrics used to evaluate the value of any investment are not a surprise given the choices, Omdia would argue that TCO is being replaced by outcome-based metrics".

– Roy Illsey, Chief Analyst, Omdia

21. When comparing pricing across cloud providers, I base my judgement on (choose one):

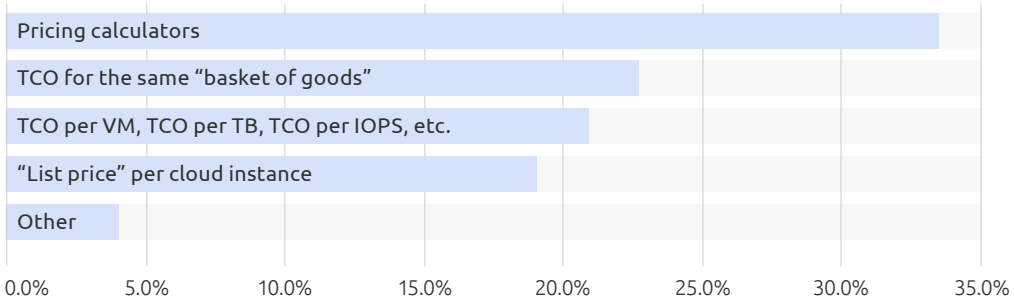

"No wonder pricing calculators come first. Whoever used public clouds before knows that estimating costs in those environments is a real challenge. All those efforts boil down to theoretically just a simple question: how much do I realistically have to pay for a given resource or set of resources. This is well reflected in the next two most popular responses too".

– Tytus Kurek, Product Manager, Canonical

22. Public cloud costs are (choose one)?

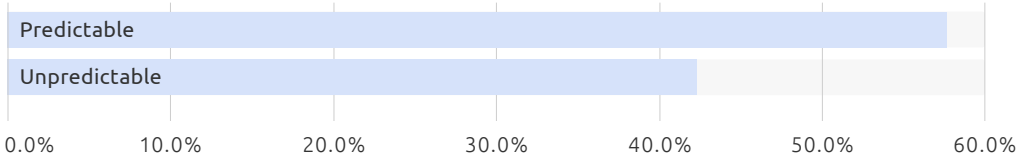

"This ties with our own research, which suggests that many organisations have so far underestimated the amount of forward planning and cost/risk management required in making the decision to move operational apps to the public cloud – and indeed, in ensuring the ability to move apps between public and private clouds as needed. Partly it's due to a relatively immature market, but a lot comes down to a lack of discipline, governance and visibility - and then there is a key mindset issue. With private cloud, capacity planning is pretty much essential on-site because you are generally working with finite resources. However, the public cloud encourages an 'abundance' mindset where you treat resources as practically unlimited. Capacity planning is therefore optional - until, of course, you realise how much a lack of modelling, planning and discipline are costing you".

– Bryan Betts, Principal Analyst, Freeform Dynamics

23. Public cloud costs can be estimated based on workload resource requirements (e.g. number of VMs, their flavours, amount of storage, required IOPS, required bandwidth, etc.) (choose one):

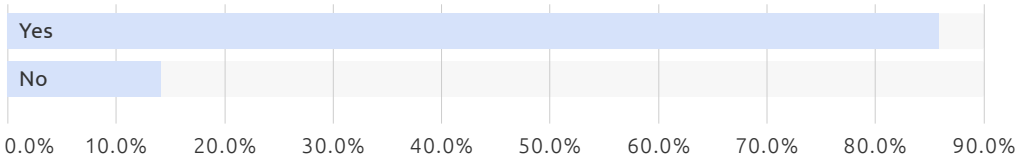

"The results are not wrong, the cost can very easily be estimated. However, it is the behaviours and processes that also need to change to ensure estimated costs are achieved. Put bluntly, this means the cloud equivalent of switching the lights off when you exit the room, otherwise the estimates will be out by a factor of more than 3".

– Roy Illsey, Chief Analyst, Omdia

"Given the nature of those answering the survey, the results are somewhat reassuring, as they imply a degree of confidence in public cloud costing among these tech and cloud-savvy respondents. However, reading it alongside the previous question is rather worrying, as it implies those cost estimates will be loose ones at best".

– Bryan Betts, Principal Analyst, Freeform Dynamics

24. Private cloud costs are (choose one)?

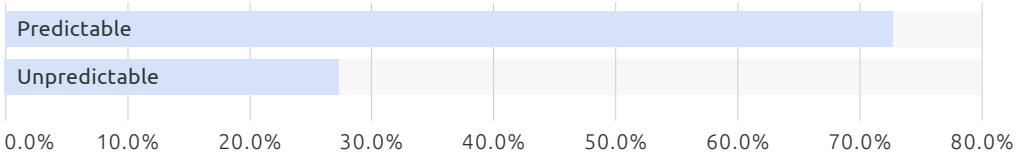

"Going on to contrast the respondents' expectations of public and private cloud, it's striking - and logical - that they believe private cloud costs are more predictable".

– Bryan Betts, Principal Analyst, Freeform Dynamics

25. Private cloud costs can be estimated based on infrastructure resource requirements (eg: hardware list price, typical hosting charges, cloud software license costs, average staff salary, etc.) (choose one):

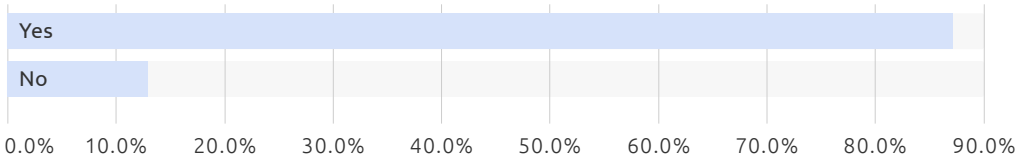

"While on-premises (private cloud) costs are more predictable as CIOs have more experience withf this type of estimations, there are still some costs that cannot be fully relied on. For example, the world has grown accustomed to low inflation, but we are now in a period of relative high inflation, so factoring in these costs is less predictable".

– Roy Illsey, Chief Analyst, Omdia

26. Hybrid cloud environments: Why would/did you build one (choose one)?

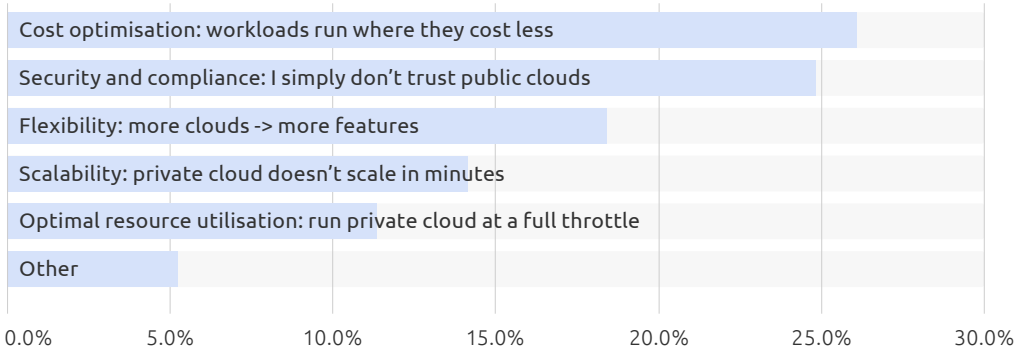

"Even though building hybrid clouds usually entails significant initial investments, clearly most of the respondents believe that there is a business case behind them. This is very much in line with Canonical's observations of the market and the cost simulations presented in this report. If it works from an economic standpoint, it's definitely worth the effort".

– Tytus Kurek, Product Manager, Canonical

27. Multi-cloud environments: Why would/did you build one (choose one)?

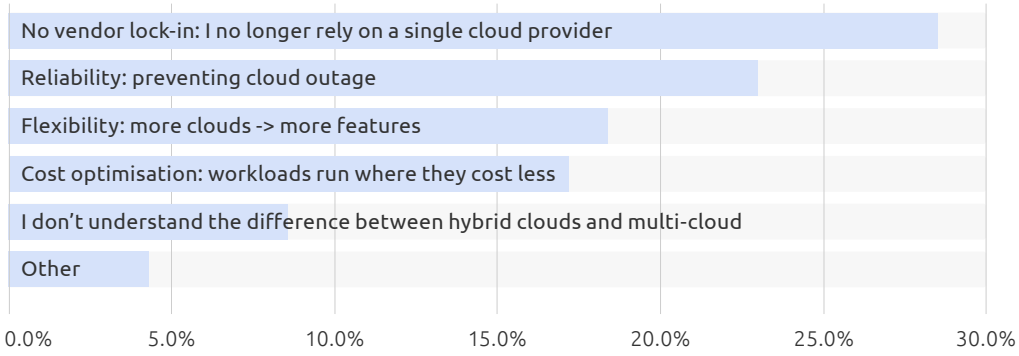

"Clearly, hybrid public/private cloud is preferred for cost optimisation, while the perceived advantages of multi public cloud are much more to do with flexibility and reliability. There's two things at play here. First, simply whether these perceptions will match the reality of the service level agreements offered by cloud providers. And second, the cost discussion is quite different for on-site vs public cloud, and the questions here don't (or can't) reflect that. On-site, yes the core resource costs are predictable, but it can be harder to allocate operational costs - plus you usually end up over-provisioning. There's also the issue of costing for infrastructure upgrades, etc".

– Bryan Betts, Principal Analyst, Freeform Dynamics

28. What is the most effective way to reduce TCO (choose one)?

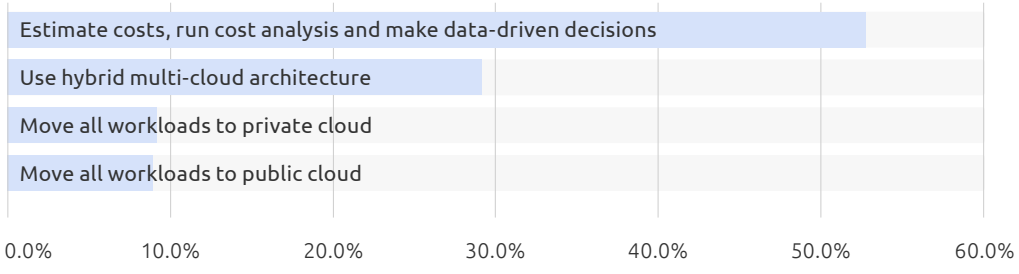

"These results are surprising to me. Almost 50% of respondents think that using one or the other cloud architecture is going to be a panaceum for their pains. Why not avoid any biases and base the decision on the data?"

– Tytus Kurek, Product Manager, Canonical

29. How do you track your personal expenses (choose one)?

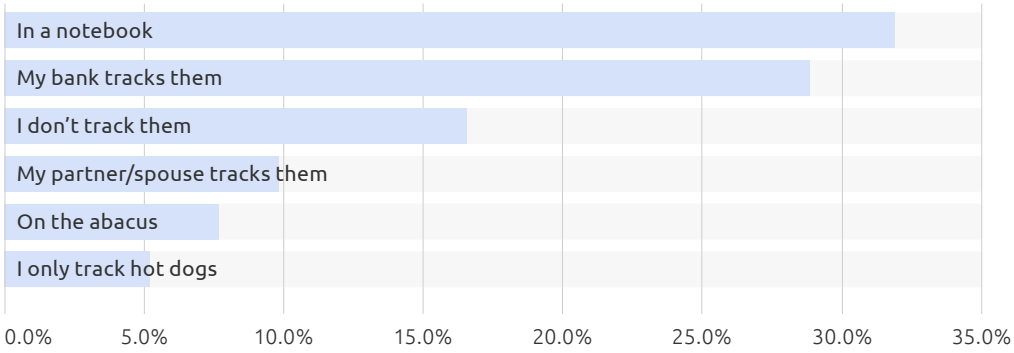

"In the past, during my university days, I used a mobile app to track every single personal expense. Later on, I analysed my expenses and tried to optimise them for the future. When I got a well-paid job, I gave up, but I still think it was a good habit."

– Tytus Kurek, Product Manager, Canonical

30. How much does your organisation spend on cloud infrastructure every year?

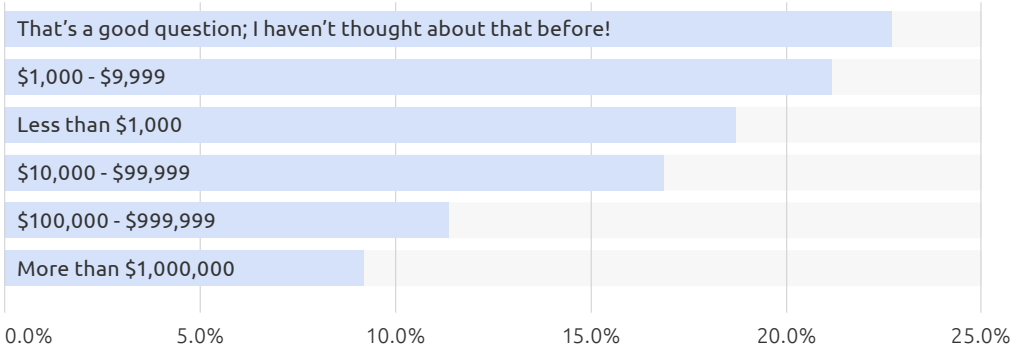

"No kidding! More than 20% of respondents don't know how much their organisations spend on cloud infrastructure every year. Well, we already know that most of them are engineers, but this is still surprising. I'd like to see how many IT managers and business executives answered this question that way. I'd advise them to go and check their bills!"

– Tytus Kurek, Product Manager, Canonical

31. Do you feel your organisation pays a lot for cloud infrastructure?

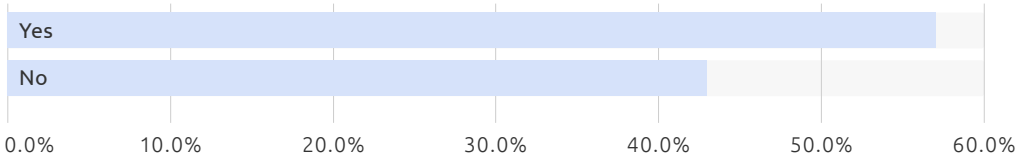

"This is surprising only for the fact it is not higher in terms of people who feel the cloud is expensive. As stated earlier, the cloud, if used correctly, is not on average any more expensive than on-premises. The value of cloud is the flexibility it provides. The issue is, it must be managed otherwise it can become expensive. This result indicates that too many organisations have not adopted FinOps as a practice".

– Roy Illsey, Chief Analyst, Omdia

32. If your business/workloads grew 10 times, how would that affect your perception of cloud infrastructure spendings?

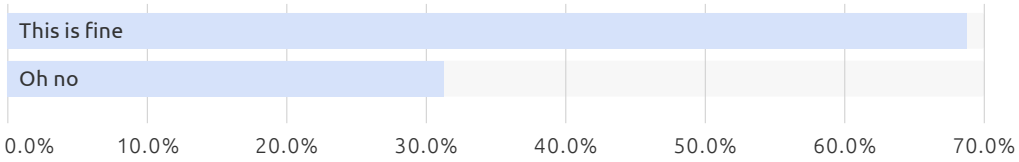

"I'm having a difficult time trying to analyse these results. In the previous question the majority of respondents said that they feel their organisation pays a lot for the cloud. But here, most of them say that it's fine if they started paying 10 times more. Is it because their organisation has the budget? Or was the question not clear enough?"

– Tytus Kurek, Product Manager, Canonical

33. Over the last two years, my organisation has seen an increase in cloud infrastructure spending:

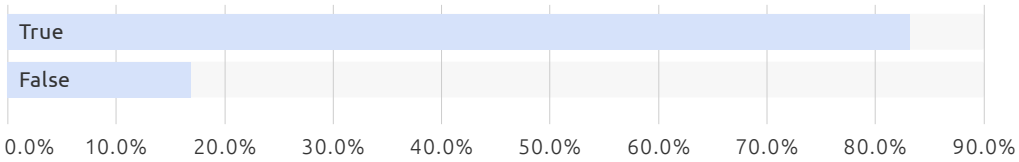

"This agrees with Omdia's research, the impact of the COVID-19 pandemic saw organisations adopt cloud to continue to offer services and operate the business. In fact, we saw workloads running the cloud jump from 25% in 2019 to 44% in 2021".

– Roy Illsey, Chief Analyst, Omdia

34. Public clouds: where do you think your organisation spends most of its budget?

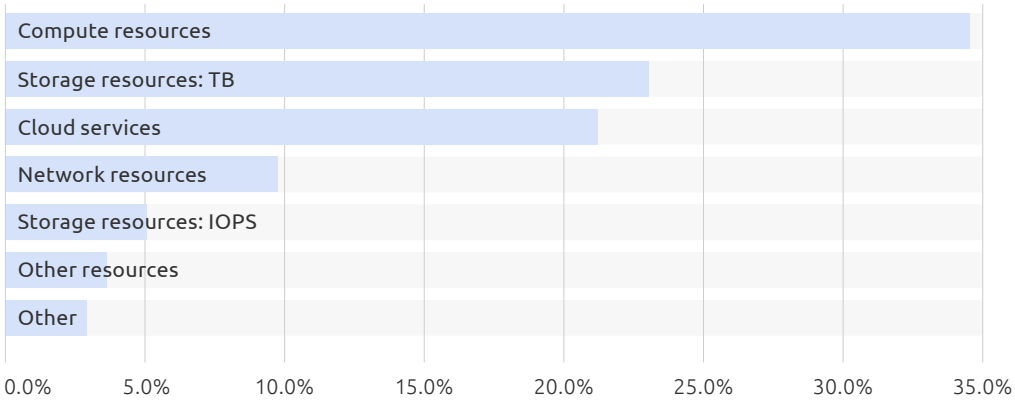

"Compute resources are undoubtedly the most consumed cloud resources these days, but people often underestimate the importance of storage. Remember that the amount of data being produced by humans keeps growing exponentially. It would be interesting to see whether this perception changes over time. I'm pretty sure there's going to be a moment when storage resources take the biggest chunk of organisations' budgets".

– Tytus Kurek, Product Manager, Canonical

35. Public clouds: do you receive any discounts from your cloud service provider?

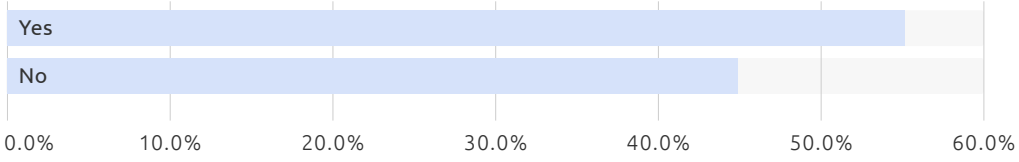

"This stresses the fact that list prices are meaningless for cost estimates and correlates with responses to question 21. This point is well-addressed in this report, however".

– Tytus Kurek, Product Manager, Canonical

36. Private clouds: where do you think your organisation spends most of its budget?

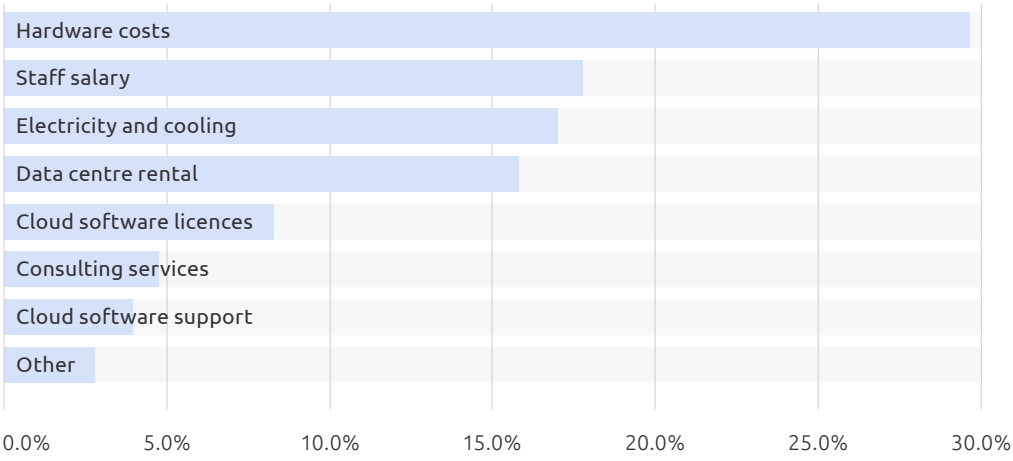

"It will be interesting to see how the perception of private cloud complexity and expense will evolve. More and more hardware and software suppliers now offer consumption-based pricing models for on-site systems - including complete packaged private cloud set-ups - often with an 'evergreen' component where the supplier handles upgrades for you".

– Bryan Betts, Principal Analyst, Freeform Dynamics

37. If not using a public cloud, why not?

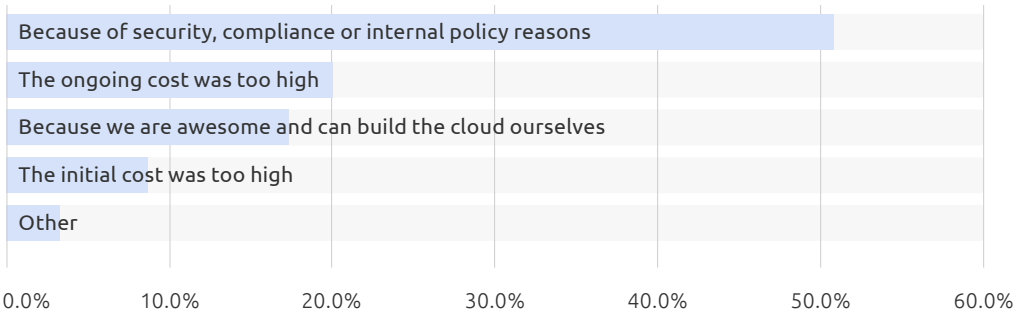

"I am a little bit sceptical about this result. Recent reports show that most attacks and data breaches occur in private clouds, not public clouds. Concerns around compliance, internal policies and data sovereignty seem to be valid, however. I'm not sure how the initial cost could be too high for public clouds".

– Tytus Kurek, Product Manager, Canonical

38. If not using a private cloud, why not?

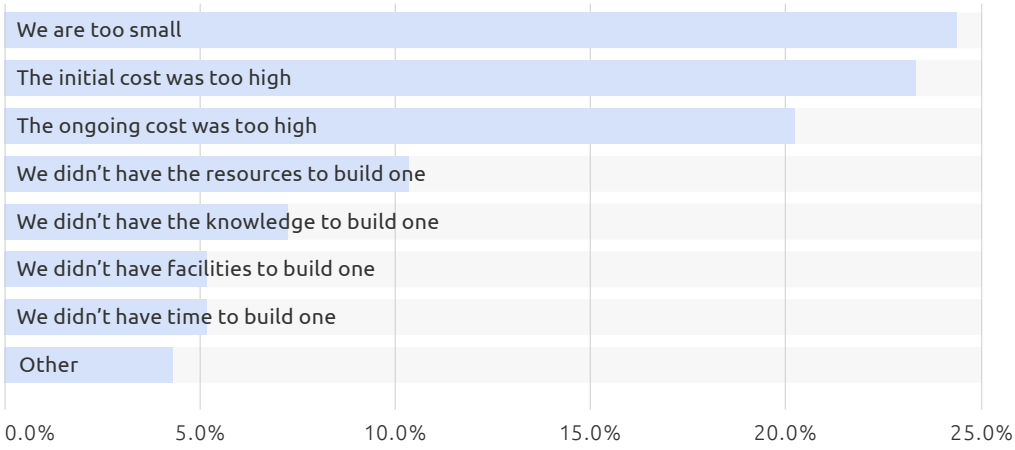

"There are multiple reasons not to build a private cloud. On the other hand, there are multiple reasons to build it. The most important thing is to be data-driven and avoid any kind of bias. Put emotions aside and analyse all pros and cons".

– Tytus Kurek, Product Manager, Canonical

39. When I tip a waiter, I:

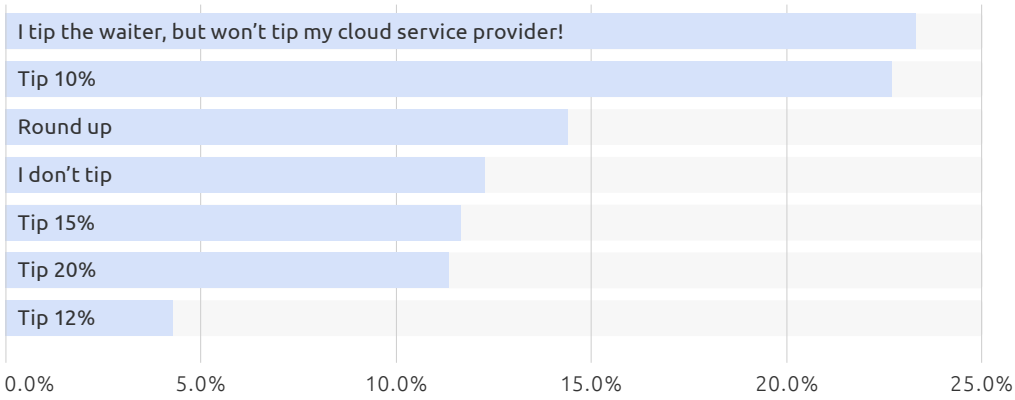

"This question has a significant cultural bias, in the UK we tend not to tip as people are paid to deliver a service. While in other countries, tipping is considered part of the cost of the service and must be added on. The interesting point is the top response: people would tip a waiter but not a CSP. Is this because of the quality of service or are most respondents thinking the service is what they are paying for. I think it is the latter."

– Roy Illsey, Chief Analyst, Omdia

Conclusions

Cloud infrastructure spend usually accounts for a significant portion of an organisation's budget. According to a survey conducted by Canonical in 2022, more than 55% of respondents feel that their organisation pays a lot for cloud infrastructure. Moreover, as businesses move towards SaaS and start using more and more cloud services, these expenses tend to increase over time. This was clearly visible during the COVID-19 pandemic when many companies switched to a remote way of working. As a result, more than 80% of survey respondents noticed an increase in cloud infrastructure spending over the last two years. Yet a lot of them don't know ***exactly how much*** they spend.

When comparing pricing across various cloud platforms, organisations look mostly at the TCO. List prices from public cloud providers and service fees from private cloud vendors are of secondary importance. There is a common understanding that there are multiple ingredients that need to be taken into account when running cost estimates. Most of them believe that both public and private cloud costs can be estimated based on resource requirements, hardware prices, staff salary, etc. TCO calculators seem to be the preferred way to run cost estimates.

More than 50% of the respondents believe that the most effective way to reduce TCO is to estimate costs, run cost analyses and make data-driven decisions. This is not surprising. The report demonstrates that cloud cost optimisation is not a trivial task; it requires an analytical approach. While public clouds are the most cost-effective option in cases that require relatively light workloads, more advanced scenarios may require the use of a hybrid multi-cloud architecture.

This again goes in line with this survey results with cost optimisation being the primary reason for building a hybrid cloud. The good news is that the majority of respondents already use more than one cloud platform, so adopting hybrid multi-cloud architecture should not be a big challenge. While many organisations resist adopting private cloud because they're too small or simply can't afford one, most of those who have done so have been able to deploy and operate it themselves. And even if that remains a challenge, such businesses can take advantage of the numerous commercial services offered in the marketplace by leading private cloud providers.

"Cloud infrastructure is now well into an age of reason, where cost management is a prime concern. More and more, organisations are optimising cost by moving to multi-cloud and hybrid cloud architectures. As the de-facto standard for open source cloud infrastructure, OpenStack is a critical part of this complex equation, and Canonical's TCO calculator shows why. It reveals significant TCO savings with OpenStack over the largest cloud providers. The Cloud Pricing Survey is another in a long line of meaningful contributions Canonical has made to the Open Infrastructure community".

– Thierry Carrez, General Manager, Open Infrastructure Foundation

Learn more

Read our business guide to hybrid multi-cloud to explore Canonical's proven approach to achieving infrastructure costs optimisation.

Visit our pages to learn more about Ubuntu Pro for AWS, Azure and GCP.

Visit our pages to learn more about Canonical OpenStack, Canonical Kubernetes and other commercial services available to enterprise customers.

Estimate the TCO of your workloads on Canonical infrastructure.

Get in touch with Canonical experts.

References

[1] https://calculator.aws/
[2] https://azure.microsoft.com/en-us/pricing/calculator/
[3] https://cloud.google.com/products/calculator
[4] https://ubuntu.com/openstack/pricing-calculator

Ubuntu