Microsoft Response to Call for Inputs: Cloud services market study
3 November 2022

1. Introduction

Microsoft appreciates the opportunity afforded by Ofcom to provide comments on the call for inputs as part of its market study on the supply of cloud services to the UK. The market study is timely. The cloud services industry today is dynamic, with growing demand and an ever-increasing number of new offerings and providers. It offers customers a new choice to decide whether to rely on their own computing infrastructure, or move to the public cloud if it offers better functionality at a more attractive price. The cloud enables organisations to reimagine their businesses, automate processes, scale capacity, and create new growth opportunities. Access to the leading-edge technologies and computing power available in the cloud is important for organisations large and small to innovate quickly, increase efficiency, accelerate productivity, and compete. Going forward, legislators, regulators and other policymakers should adopt policies designed to ensure that UK enterprises, public sector organisations, and non-profit entities have access to and can choose the best cloud computing technology.

Microsoft's response to Ofcom's Call for Inputs (CFI) into its cloud services market study is structured as follows:

- Section 2 outlines Microsoft's view on the current competitive landscape and future developments in cloud services (CFI questions 4.3 and 4.5)

- Section 3 outlines Microsoft's comments on the scope of the cloud computing market study with a focus on the relevance of ecosystem competition based on the underlying features of cloud services offerings (CFI questions 4.1 and 4.5 to 4.8)

To inform its market study, Microsoft understands that Ofcom expects to consider various prior international studies. (CFI paras. 2.19 to 2.24). Although these studies can help inform Ofcom's work, the various conclusions and determinations from those studies may not necessarily reflect today's reality or the conditions that exist in the UK. As Ofcom examines important topics that include the degree of competition on various markets, ease of switching, lock-in, and scale impacts among others, it will be critical to view those issues afresh in the UK context (and not to start with a premise based on conventional wisdom from earlier studies). There is strong

competition between cloud services providers today. And that competition continues to drive digital transformation, efficiency, productivity, and innovation throughout the UK economy as well as public and non-profit sectors. Intervention to avoid future potential concerns, therefore, should be pursued with caution so as not to undermine innovation, competitiveness, and growth in the UK going forward.

2. Landscape of cloud computing in the UK

a. Growth and innovation in cloud services offerings

As the CFI recognises, cloud services are growing rapidly with constant innovation and overall lowering of prices over time.[1] For example, Gartner estimated that in the IaaS segment alone, worldwide growth was over 40 percent in 2020. Amazon Web Services ("AWS") remains the worldwide leader with roughly 40 percent share depending on which services are considered.[2] Along with AWS, Microsoft Azure, and Google Cloud Platform (GCP) are currently the largest suppliers of cloud services. Each invest billions of pounds each year in building new datacenters, expanding networking infrastructure and developing new cloud services.

The CFI classifies cloud services by service models: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). As discussed below, while these categories can be helpful conceptual tools to analyse competition in cloud services, they are an oversimplification and do not necessarily reflect how suppliers or customers approach solutions in the cloud. Moreover, there is a danger in concluding, as the CFI does, that "*each layer is notionally built on top of the previous one(s)*." (CFI para. 3.5) These models simply describe the division of responsibility or control between the customer and the cloud services provider — a dynamic that the customer controls by choosing which services it will rely upon and picking and choosing among them whether they are notionally IaaS or PaaS. A SaaS solution must be built on server hardware and be supported by platform software, but those may be assembled in a highly integrated and optimised architecture not amenable to disaggregation or switching out of the underlying "layers." Moreover, the customer of a SaaS solution, by definition, has no access to

[1] *See* https://appdevelopermagazine.com/why-the-cost-of-cloud-computing-is-dropping-dramatically/; https://journal.uptimeinstitute.com/cloud-generations-drive-down-prices/.

[2] In 2020, Gartner estimated that Amazon had over 40 percent share of the revenue for IaaS. https://www.gartner.com/en/newsroom/press-releases/2021-06-28-gartner-says-worldwide-iaas-public-cloudservices-market-grew-40-7-percent-in-2020.

the underlying infrastructure or platform software; he or she only experiences the finished application.

In analysing competition in cloud services, looking at IaaS and PaaS solutions on the one hand and SaaS solutions on the other can be helpful (CFI para. 3.11) primarily because there are "important differences in the features of the services and types of suppliers" as the CFI notes in para. 3.11. Indeed, IaaS and PaaS are primarily sold on a consumption (or pay-as-you-go) basis and used by developers to create their own cloud-based apps and solution for their own internal or external use. "Public cloud" providers typically provide services that span IaaS and PaaS. There are, however, also numerous third-party providers of cloud solutions that offer their own services that they make available on different public clouds for use by developers to create their own cloud powered solutions.

In contrast, cloud services that fall into the SaaS category, whether offered by the provider of the underlying hardware infrastructure or built on infrastructure controlled by a third-party, are primarily offered on a subscription basis per user. They are used by consumers or organisations and provide complete ready-to-use application functionality. The differences in business models, use cases, and customers between IaaS / PaaS and SaaS also means that even where a provider offers all services in both categories, the customer demand and corresponding sales motion is entirely different between IaaS/Paas and SaaS.

At the same time, these categorisations are merely tools to help understand the types of cloud services that are developed and should not themselves be considered as relevant markets. Nor do they enable consistent delineations between different services. Customers look to the public cloud to complete certain tasks or workloads, using the combination of cloud services that best meet their needs. This might not necessarily be an "IaaS" offering or "PaaS" offering. In fact, for many services the distinctions will be blurred with no meaningful and consistent way to distinguish between IaaS or PaaS. For example, Azure customers commonly provision a virtual machine (VM) with a certain amount of memory, storage, and compute power running either Linux or Windows Server. This is often a single stock keeping unit (SKU), but it is not totally clear whether the provisioning of a Windows VM falls in the IaaS category, because the user is buying compute power, or falls in the PaaS category, because the SKU includes an operating system in the VM. Similarly, storage is generally considered to be IaaS, but there are many solutions for managing the storage of data, such as a data warehouse offer from a company like Snowflake.[3]

---

[3] https://www.snowflake.com/data-cloud-glossary/data-warehousing/.

Snowflake data cloud offers a SQL data warehouse, which would likely be considered in the PaaS category. In sum, the distinction between IaaS and PaaS is blurry and customers look to the public cloud to find the best solution to address their needs irrespective of strict categorisation.

b. Providers of cloud computing services

As is recognised by the CFI, AWS, Microsoft Azure, and GCP are currently leaders in the provision of public cloud services. That said, the needs and demands of cloud customers vary widely. We do not expect that any one or even three public clouds will be the only entities to meet these needs in the future. In fact, other firms are also competing, innovating, and growing today or are poised for entry. For example, Oracle is aggressively growing its infrastructure cloud business, competing for customers worldwide.[4] IBM too offers a full stack of public cloud services with a focus on enabling hybrid cloud deployments and AI services.[5] And there are a variety of relatively smaller but successful cloud providers. For instance, Digital Ocean, with 14 datacenters around the world, distinguishes itself from other providers by easing the onboarding processes, simplifying pricing, providing better support, and offering a curated set of services that target start-ups and SMBs.[6] And, of course, there are thousands of providers of different cloud services that are offered as hosted solutions on infrastructure in datacenters offered by third-parties such as Amazon, Microsoft, Google, IBM, and others.

Because cloud services generally benefit from economies of scale, Microsoft acknowledges that entering from scratch as a new public cloud provider can be challenging. This does not mean, however, that entry is unachievable or that the actual or potential entry is not an important competitive constraint. With respect to entry and potential competition, there are many different potential entrants. As acknowledged by the CFI at para 3.16, there are at scale public cloud services offered by Alibaba, Tencent and Huawei which primarily operate outside of the UK.[7] They are well poised to expand into the UK in light of their existing global network of datacenters and their expansive and mature IaaS and PaaS offerings. Another potential scenario for entry is by large well-resourced organisations, which have an incentive to develop an at-scale cloud for

[4] *See, e.g* https://siliconangle.com/2022/09/12/strong-cloud-growth-propels-oracle-results/.

[5] *See* https://www.ibm.com/cloud/products.

[6] *See* Q2 2022 Earnings Presentation for Digital Ocean (details growth opportunity and differentiated offering) (available at https://s27.q4cdn.com/619704647/files/doc_financials/2022/q2/Q2'22-Earnings-PresentationFINAL-(1).pdf).

[7] Alibaba's AliCloud includes a datacenter in London, UK, since 2018. *See* https://www.alibabacloud.com/globallocations#:~:text=Data%20Centers%20Around%20the%20World%201%20Asia%20Pacific,1%20%28Tokyo%29%20 Availability%20Zone%3A%202%20...%20More%20items.

their own use and then open it up to deliver services to others. There are many well-resourced providers who fall into this category, e.g., Facebook, Salesforce, and Apple that all run

their own massive datacenters to power their popular offerings. As the CFI recognises,[8] the scale of their own operations would enable entry into the IaaS and PaaS segment. This has already been proven by the entry and expansion of Amazon when it leveraged the data centers that powered Amazon's retail marketplace to first introduce AWS.

In addition, the ability to run software on-premises or develop a private cloud exercises a competitive constraint on public cloud providers. Customers need not move their computing workloads to the public cloud. They have choice and complete control. Moreover, firms can and do choose to use both internal and external solutions and shift workloads and/or solutions between them. For example, Walmart has developed a private cloud and at the same time uses public cloud services from GCP and Microsoft Azure.[9] Customers will use public cloud services only if they are more attractive than their on-premises or private cloud solutions. Since the main customers of the public cloud are sophisticated developers, they can also create their IT solutions with a potential to be able to switch or with contractual guarantees that ensure such a transition remains an economically viable option in the future.

c. Factors that customers consider in choosing public cloud services

The CFI speculates as to the features that customers consider when choosing cloud services. (CFI para. 4.15) While those listed are undoubtedly important to some customers, the needs and demands of cloud customers vary widely. The customers range from small and medium businesses to large enterprises, public sector agencies, non-profit organisations, educational or research institutions. Different types of organisations will face different constraints including their respective budget and financial resources, procurement processes, geographic footprints (local, regional, or international), legacy IT infrastructure, and access to skills and capabilities. In addition, customers come from different industries and this also impacts what specific cloud

[8] The study explains that it will "*explore the extent to which SaaS-only or private cloud-only providers may provide an entry constraint on existing public cloud providers, as well as any scope for potential competition concerns in public cloud to have an impact on either of these segments.*" Call for Inputs at para 4.8.

[9] https://www.techtarget.com/searchcloudcomputing/news/252522631/Walmarts-multi-cloud-strategy-cutsmillions-in-IT-costs.

services they need as well as dictates requirements in the form of minimum compliance, privacy and security standards. In short, it is challenging to generalise with certainty which factors matter most to customers when choosing cloud providers.

Of course, one key consideration for any cloud customer is the total cost of the solution, including how it compares to the costs associated with other options such as relying upon their own hardware infrastructure. In our experience in addition to price, many customers take into consideration at least the following additional factors:

- **Scalability.** Scalability refers to the ability of a system to handle increased load, i.e., can a cloud provider scale services automatically to match demand, ensuring capacity during workload peaks and returning to normal automatically when the peak drops. Different customers will have different needs for scalability.[10]

- **Resiliency.** Resiliency is a system's ability to recover from failures and continue to function. It's not only about avoiding failures but also involves responding to failures in a way that minimizes downtime or data loss. Different customers will have different needs for resiliency.[11]

- **Sustainability.** Cloud migration can help customers meet sustainability goals. As pressure builds for greater progress on such goals, we expect that customers will increasingly consider the environmental impact of different cloud providers when choosing among them.[12]

- **Compliance.** Many customers need cloud services providers that comply with national, regional, and industry-specific requirements governing the collection and use of data.[13]

[10] For more information *see* https://learn.microsoft.com/en-us/azure/architecture/framework/scalability/designscale.

[11] For more information *see* https://learn.microsoft.com/en-us/azure/availability-zones/overview.

[12] For more information *see* https://www.microsoft.com/en-us/sustainability/azure.

[13] For more information *see* https://learn.microsoft.com/en-us/azure/compliance/ and https://azure.microsoft.com/en-us/explore/trusted-cloud/.

- **Security.** Cybersecurity is a paramount concern for almost all cloud services customers.[14]

- **Privacy.** Many customers want to be able to control their data and where it is located as well as ensure that their cloud services provider meets applicable privacy laws and standards such as GDPR. For some customers the availability of independent audits may be relevant.[15]

3. Cloud services are not protected from competition by "ecosystems" as is sometimes the case with consumer platforms.

The CFI details its plan to "*assess the extent to which hyperscalers are competing on the basis of their ecosystems and how this might shape competition in the future*." (CFI para. 1.23) The CFI explains that ecosystems can be "*understood as a collection of complementary products and services that work together to create utility for customers.*" (CFI para. 4.23) More specifically, the CFI asserts that cloud services may amount to ecosystems, because of the "*full-stack service offerings that each of the hyperscalers provides*".[16] (CFI para. 4.24) The CFI appears to assume that because cloud providers offer access to hardware resources and software building blocks for cloud computing solutions to make use of the hardware, there is an ecosystem. But before making such an assumption, it is necessary to consider the concept of an ecosystem from an economic point of view and then assess if the factual circumstances under consideration fit an accepted economic definition. An ecosystem can be described as a large set of products that are highly complementary. In addition, ecosystems are typically characterised by large relative

[14] For more information *see id.*

[15] For more information *see id.*

[16] The CFI cites as an example the CMA study that properly identifies an ecosystem, explaining that *"in the mobile ecosystems market study, the CMA identified several gateways between users and businesses – mobile devices and operating systems, app stores, and mobile browsers and browser engines – that Apple and Google control."* This type of ecosystem and the concerns raised by them are very different than the development of ever richer and richer functionality that clouds make available to developers.

switching costs if a customer wants to switch away from one product to another. These characteristics do not exist in the public cloud.

a. Switching and multi-homing behaviour between cloud services indicate there is no ecosystem

It is true that public cloud providers compete to provide ever richer functionality to enable cloud customers to accomplish more at lower cost and with less effort. This means that customers can find a vast array of complementary functionality to develop their IT solutions at each public cloud provider. However, Ofcom should be careful not to draw any undue analogies from the consumer space, e.g. mobile operating systems, which operate very differently from cloud services.

Ecosystem competition may be a major issue with consumer platforms, where end users find value by ensuring that all of the apps and services they want to use are available on a particular

platform. Users do not multi-home, for example, by using one mobile phone for calling and photography and another mobile phone for texting and personal finance apps. Similarly, the mobile user will be interested in having multiple devices that all work together as part of the platform, such as a mobile phone, connected speakers, car operating systems, wearable devices and the like. This can give rise to ecosystem effects because devices, operating systems, and apps are not only platforms, but changing any complementary component may imply that expensive devices have to be changed or customers have to learn how a different operating system works. In such cases the existence of large switching costs is much more plausible. A customer will not buy a new car in order to switch mobile phones and, instead, if access to the phone in the car is important, they will stick with the existing compatible option.

This does not hold for cloud services. Switching costs between cloud providers depend on the nature of the solution and are largely in the customer's control. Unlike consumer products, the customers of cloud services are typically sophisticated organisations with developers with specialised knowledge making purchase decisions, which enables the customer to make well-informed choices with a clear understanding of the trade-offs involved in optimising workloads to reduce the costs of migration.[17] This may involve customers obtaining cloud solutions from multiple suppliers, and using technical solutions that allow mixing-and-matching between the

[17] *See, e.g.,* "Managing Technical Lock-in in the Cloud," (17 Dec. 2019) (available at https://www.gov.uk/guidance/managing-technical-lock-in-in-the-cloud).

different suppliers' offerings. For example, Walmart recently shared its experience building a cloud native platform for its ecommerce retail website. In doing so, it leveraged Kubernetes (container technology) to enable its platform to run across two public clouds – Microsoft Azure and GCP – and to take advantage of its own OpenStack-based cloud infrastructure.[18]

Open-source technologies, like Kubernetes containers, are prevalent across public clouds and give customers – like Walmart and others – confidence that if they build their solutions using such services, they will be able to migrate their workloads to other public clouds. In addition, commercial first- and third-party PaaS services have emerged that facilitate multi-cloud usage and migration. As an illustration, Microsoft created Azure Arc to enable customers to manage, govern, and secure workloads deployed across on-premises, multi-cloud, and in edge environments from a single place. Another example is Snowflake, a popular third-party cloud computing service, which offers a data platform on all major clouds, including AWS, Azure, and

GCP. Customers using Snowflake can easily move and access their data across those clouds.[19] Snowflake explains that it "*ported [its] software by building a cloud-agnostic layer, abstracting the specificity of the underlying cloud infrastructure. This means any application running on Snowflake Cloud Data Platform is also cloud agnostic. This aspect is very important since avoiding any cloud lock-in is one of the key benefits of using Snowflake*."[20] Indeed, many providers are offering solutions across multiple clouds to enable better multi-cloud strategies.[21] These solutions show that in contrast to an ecosystem that is characterized by significant switching costs and barriers to interoperability, there are interoperability solutions that support substitution by shifting services between clouds and switching to new solutions, whether that means switching an existing workload or selecting a different provider for future cloud based workloads. Indeed, for each incremental workload or solution, customers can choose which cloud provider to use (or

[18] *See* https://www.techtarget.com/searchcloudcomputing/news/252522631/Walmarts-multi-cloud-strategycuts-millions-in-IT-costs.

[19] Of course there are many other ways to move data across clouds as well. Use of a common database structure across multiple clouds can be helpful, but could also be done with solutions like MySQL or other well understood solutions.

[20] https://www.snowflake.com/blog/how-snowflake-delivers-a-single-data-experience-across-multiple-clouds-andregions/

[21] *See* for example the Enterprise Monitoring Tools from Datadog that enable monitoring infrastructure across clouds and on premises and that works across many major clouds. https://www.datadoghq.com/enterprisemonitoring-tools/.

whether to deploy on-premises or in a private cloud). Cloud computing customers are not confined to using one cloud provider for all their needs, and they do not.

The Flexera 2022 State of the Cloud Report confirms that multi-cloud is "the de facto standard" and reported that "79 percent of respondents said they're incorporating multiple public clouds, and 60 percent report using more than one private cloud—an increase from last year."[22] As the figure below shows, customers choose to use multiple cloud providers for a variety of scenarios.

FIGURE 11

Apps siloed on different clouds remains the most common implementation, but *DR/Failover* between clouds saw a significant increase YoY.

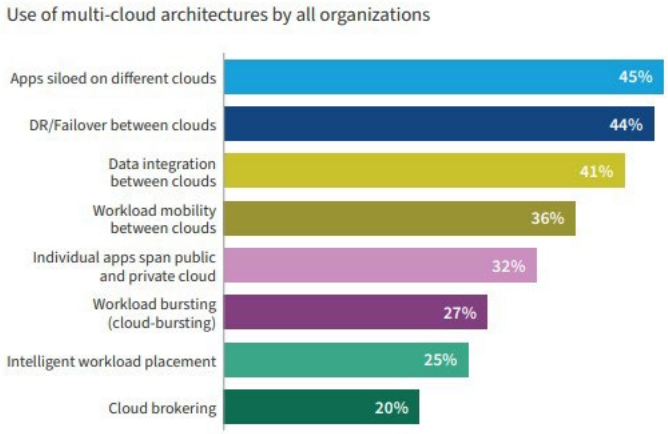

[22] *See Flexera 2022 State of the Cloud Report* at pg.17 (available at https://www.flexera.com/?elqTrackId=d2174077cc874550a3d7e8bd2fc403bf&elqaid=6546&elqat=2).

The most common multi-cloud use is apps siloed on different clouds. In Microsoft's view, this does not lead to the conclusion in the CFI that "*[t]his type of multi-cloud use is less likely to facilitate switching.*" (CFI para. 3.37) The fact that customers run different apps on different clouds says nothing about whether they can be easily switched from one cloud to another. In fact, as the survey results suggest, customers do deploy the same workload across clouds for failover (44% of customers) and integrate data between different clouds (41%) which does imply the use of different clouds for similar workloads. These two categories are in addition to workload mobility between clouds (36% of customers). Indeed, the APIs made available by cloud vendors are made broadly and openly available to be called by applications and even applications running in different clouds. For example, Microsoft documents all of its Azure APIs on Github and developers can call them within Azure, from on-premises or between clouds.[23]

b. Cloud services do not involve meaningful network effects like are typical with ecosystems

An important and related point is that cloud services do not exhibit any meaningful "network effects". Network effects occur when the value of a product, service or platform depends on the number of buyers, sellers or users who use it. Typically, the greater the number of users, the greater the network effect. "Direct" network effects occur when increased usage benefits all users on one side of a platform. "Indirect" network effects exist where the value of the platform

---

to one group depends on how many members of another group participate. In other words, the value of the services that a two-sided platform provides increases as the number of participants on both sides of the platform increases. For example, a credit card is more valuable to cardholders when more merchants accept it and is more valuable to merchants when more cardholders use it.

These network effects can also occur in certain IT settings – for example, consumer uptake of applications written for a desktop operating system (such as word processors or spreadsheets) can make the platform more valuable to all users because they can exchange data and information with an expanding circle of users (direct network effect). In addition, more application users can make the platform more valuable to software developers, creating a virtuous circle of usage and development (indirect network effect).

---

[23] *See* https://github.com/orgs/Azure/repositories?type=all.

In the case of cloud services, however, these network effects are largely absent. Companies choosing a cloud services provider are indifferent as to whether its public cloud has more or fewer end users. At the same time, a public cloud does not generate incremental additional value to other users from hosting many other users. While having many customers can create economies of scale for the cloud service provider, they provide no direct benefit to other users. Stated differently, additional business is additional business for the cloud service provider, but the whole is not greater than the sum of its parts. As a result, a would-be cloud services entrant can start offering cloud services without facing a network effects barrier to entry on the user side.

Similarly, the fact that some third parties elect to provide PaaS services on some public clouds does not create significant indirect network effects as would exist for a PC or mobile operating system. For a PC or mobile operating system, end users must ensure there are a sufficient number of applications written and available for those operating systems such that they meet all their needs. An operating system without the requisite applications cannot be competitive because users are unlikely to purchase an additional device to get access to a missing application. Cloud services, however, are a more fundamental input. Customers pick and choose solutions and build compelling applications and services using the cloud infrastructure services of their choosing. Even if a third-party solution that the customer wants for a particular workload is not available on a particular cloud, it does not prevent the customer from deploying a different workload that do not rely upon the third-party solution on that cloud.

Many cloud services are designed to enable usage across multiple different cloud environments. And, even where a third-party cloud service has designed itself to optimise and take advantage of specific cloud infrastructure, they are likely to do so for all major cloud providers (and not simply one “ecosystem”). There is relatively little preventing such ISVs from supporting additional cloud infrastructure offerings. In many cases, similar functionality to that provided by the thirdparty solution is already provided by (or could be provided by) the cloud provider directly, enabling customer choice between clouds even when third-party solutions are not offered on a less popular cloud.

As such, although public clouds are often referred to as a “platform,” this can be considered as true only from a technological perspective: developers build applications and deploy workloads on the datacenter platform offered by the cloud services provider. However, from an economic perspective, public clouds are primarily a *one-sided* factor of production. They serve one set of users: customers developing solutions for their own internal or external use. Developers consume the services and pay directly for their use based on consumption. Thus, the public cloud

is best thought of as an input that does not generate the same concerns regarding network effects and ecosystems that are so prevalent for consumer-oriented two-sided digital platforms. This has implications for the economic incentives and competition in the markets.

c. Self-preferencing risk is low in cloud computing services

Because the CFI focuses on concerns associated with consumer digital platforms, including concerns about network effects and ecosystems, it is not surprising that it also lists selfpreferencing as a concern to explore. In particular, the CFI explains that there could "*be hyperscaler self-preferencing where it directs its customers to its own PaaS services (rather than ISV services).*" (CFI para. 4.33) Unlike consumers, however, customers of cloud services are sophisticated developers where concepts like default bias and nudging towards the platforms' preferred services do not drive usage. Rather, developers determine and choose the solutions that work best with what they are building or deploying. The fact that cloud services are typically consumed by sophisticated developer customers means and the prevalence of cross cloud third party products also reduces a likelihood of self-preferencing. Efforts to interfere with use of the products and services of the developers' choosing is likely to do little more than irritate the developer. Thus, carrying over the concerns, like self-preferencing, that arise in the context of consumer digital platforms to the public cloud could result in very odd outcomes.

4. Conclusion

1. Microsoft looks forward to cooperating with Ofcom on its market study and in particular on the important issues relating to cloud computing outlined above. The benefits and flexibility brought about by the cloud have helped transform the technology sector. Ensuring a vibrant and competitive cloud industry is critical to the future of the digital sector in the UK and the world.