# Cloud services market study

## Final report

**Redacted [✂] for publication**

Published 5 October 2023

# Contents

## Section

# Overview

1.1 Cloud computing is being rapidly adopted by businesses across the economy and has become an essential part of how digital services are delivered to consumers, including in the telecoms and broadcasting sectors. Ofcom has carried out a market study into the supply of cloud services in the UK to explore if these markets are working well and whether any regulatory action is required.[1] This final report sets out our findings and recommendations.

## Our final report – in brief

'Cloud computing' is the provision of remote access to computing resources (such as compute, storage and networking) on demand and over a network. Cloud computing has both transformed the way businesses and organisations of all types and sizes run their operations and become a critical input to the digital services we all rely on each day.

Our study is focused on 'cloud infrastructure services', which are built on physical servers and virtual machines hosted in data centres around the world. Cloud infrastructure provides the foundation for how software applications are developed and run. This consists of products called infrastructure as a service (IaaS) which includes storage, computing and networking, and platform as a service (PaaS) which includes the software tools needed to build and run applications. The market for cloud infrastructure in the UK was worth £7.0 billion to £7.5 billion in 2022.

There are two leading providers of cloud infrastructure services in the UK: Amazon Web Services (AWS) and Microsoft, who had a combined market share of 70% to 80% in 2022.[2] Google is their closest competitor with a share of 5% to 10%. Collectively these firms are referred to as the 'hyperscalers' and the vast majority of customers use their cloud services in some form. A diverse set of independent software vendors (ISVs) build their products on cloud infrastructure from the hyperscalers, but also compete directly with some of their services.

Our study has found that competition between cloud providers is mainly focused on attracting new customers when they first move into the cloud. We see evidence of some positive outcomes for customers, including product innovation, discounts and a wide choice of software services from ISVs. However, our view is that competition is being limited by market features that make it more difficult for customers to switch and use multiple suppliers (known as 'multi-cloud'). The features we are most concerned about are:

- egress fees are the charges that customers pay to transfer their data out of a cloud. The cost of transferring data between rival providers can discourage customers from using more than one cloud provider and in some cases make switching more costly.
- technical barriers mean that customers need to put additional effort into reconfiguring their data and applications to work on different clouds. A lack of interoperability and portability can restrict the ability of customers to switch and multi-cloud.

[1] On 6 October 2022, we published a market study notice in accordance with section 130A of the Enterprise Act as amended and applied by section 370 of the Communications Act.

[2] We used a combination of data sources to estimate market shares and present our estimates in ranges for confidentiality reasons. See Annex 1 for more details.

- committed spend discounts can benefit customers by reducing their costs, but the way these discounts are structured can incentivise customers to use a single cloud provider for all or most of their cloud needs. This can make it less attractive to use rival providers as part of a multi-cloud strategy.

As a result, we are concerned that a material number of customers, especially those with more complex requirements, may face significant barriers to switching and multi-cloud. We expect this will be true of an increasing number of customers as the market matures. Some customers have told us they are already concerned about being 'locked in' to their current provider.

Limits on the ability of customers to credibly threaten to switch away can reduce the competitive pressure on the market leaders, giving them a degree of market power. This creates the risk of harm for cloud customers, either by paying higher prices than would have been the case or being denied access to innovative products, which in turn can lead to negative impacts for UK consumers. High levels of profitability for the market leaders AWS and Microsoft and a gradual increase in market concentration are consistent with limits to the overall level of competition.

Looking ahead, if customers have difficulty switching and using multiple providers, it could make it harder for competitors to gain scale and challenge AWS and Microsoft effectively for the business of new and existing customers. There could be long lasting impacts if this leads the market to become more concentrated, with barriers to switching and multi-cloud allowing the market leaders to entrench their positions and avoid competing vigorously. This could have implications for ISVs, especially where they become more dependent on the market leaders for access to customers.

A cloud infrastructure market that is working well is critical for businesses across the economy and everyone who makes use of digital services. Given the concerns we have identified, we have decided to refer the cloud infrastructure market to the Competition and Markets Authority (CMA) to carry out a market investigation. The CMA will now conduct an independent investigation to decide whether there is an adverse effect on competition, and if so, whether it should take action or recommend others to take action.

## Cloud computing is important to the markets Ofcom regulates and the wider economy

1.2 Cloud computing has been widely adopted by UK businesses across the economy. Compared to the traditional model, where businesses purchase and maintain their own physical computing resources and software, cloud computing is faster to deploy, more flexible and potentially cheaper. This supports innovation and growth, for example by allowing businesses offering digital services to scale up quickly and cost effectively.

1.3 It is an increasingly important input to the different elements that make up the internet, which means it is essential for providing online services used by many UK consumers including social media, streaming, and communications services. Cloud computing is expected to underpin the development of artificial intelligence (AI) as it provides the computing resources and infrastructure needed to train and deploy AI models at scale.[3] AI is also expected to enhance the functionality of software applications that run in the cloud.

[3] See more generally, the recently published initial report by the CMA on AI Foundational Models [accessed 18 September 2023].

1.4 This technology is also changing how services in the telecoms and broadcasting sectors are being produced and delivered to consumers. In broadcasting we already see extensive use of the cloud by public service and commercial broadcasters, including growing use in the production of TV and video content. Cloud computing is expected to play an increasing role in the delivery of fixed and mobile telecoms, with partnerships emerging between cloud providers and telecoms providers in the UK and internationally.

1.5 If the markets for cloud services are not working well, there could be negative impacts for the businesses that rely on them through higher prices, lower service quality and reduced innovation, that would ultimately be passed on to UK consumers.

## AWS and Microsoft are the clear leaders in cloud infrastructure

1.6 The supply of cloud infrastructure in the UK is concentrated, especially at the infrastructure as a service (IaaS) layer, where Amazon Web Service (AWS) and Microsoft are the clear market leaders. AWS and Microsoft account for 70% to 80% of UK IaaS and platform as a service (PaaS) revenues.

1.7 AWS was first to launch cloud services in 2006 and has been able to maintain a significant share as other providers have entered the market. Our analysis indicates that AWS's profitability has been consistently high, with returns significantly above our estimate of the weighted average cost of capital (WACC) since at least 2014. Microsoft is the closest competitor and has grown its share significantly since it entered the market in 2010. We estimate that Microsoft's public cloud division, Azure, is becoming increasingly profitable and that its returns are also above our estimate of the WACC.

1.8 Google is the main challenger to AWS and Microsoft. Google entered the market in 2011 and while its share has grown in recent years, Google remains significantly smaller than the two market leaders, with a 5% to 10% UK share across IaaS and PaaS combined. Google's cloud division recently made a profit for the first time, although this was relatively low compared to the profits of AWS and Microsoft.

1.9 The hyperscalers offer a broad range of complementary services across the different layers of the cloud stack. In addition to selling their own products, they also host PaaS and software as a service (SaaS) products developed by independent software vendors (ISVs) and act as channels for customers to purchase these services, including through marketplaces. These developments suggest that AWS, Microsoft and Google are each building their own 'ecosystems', that provide customers with access to a broad portfolio of their own and others' products in a single place that work together seamlessly.

1.10 Beyond the hyperscalers, there is a range of relatively smaller cloud providers present in the UK, including some who also operate across all parts of the cloud stack. These include large technology companies such as Oracle and IBM, who both have considerably smaller market shares at around 0% to 5% of UK IaaS and PaaS revenues. These providers are more distant competitors to the hyperscalers, partly because of the difficulty of building a rival ecosystem of products delivered over a global network of data centres.

1.11 A wide range of ISVs compete mainly in PaaS and tend to specialise in a particular area, such as databases or analytics, rather than across several different product categories.

Collectively they account for a significant share of 30% to 40% of UK PaaS revenues, but our analysis suggests no single ISV has a share greater than 5%.

## Competition is currently focused on attracting new customers who are moving to cloud for the first time

1.12 The UK cloud infrastructure market is growing, with overall revenues increasing at a rate of 35% to 40% annually in recent years. It features a diverse range of customers from different sectors across the economy, each with different requirements. Some have more recently moved to the cloud, either as new start-ups or later adopters. Other more established businesses expect to move more of their data and applications into the cloud over time. Large enterprises account for a high proportion of providers' revenues and their behaviour is particularly important for the competitive dynamics of the market.

1.13 The initial choice of cloud provider is a critical moment for customers. Once a customer chooses a provider they are likely to increase their usage with that provider over time, particularly where it becomes costly to switch away or introduce an additional provider. This means competition between the hyperscalers is mainly focused on attracting new customers into their ecosystems when they first move into the cloud. Significant discounts are offered in return for committed spend by larger customers, alongside technical support to help businesses move applications into the cloud.

1.14 Once customers are established in the cloud there are clear benefits to adopting a multi-cloud strategy to get access to the best quality services, build resilience into their cloud architecture and strengthen the bargaining position with their provider. We are aware of some larger and more sophisticated customers who are adding a second cloud provider for specific use-cases. However, we have found few cases where customers are able to take an approach to multi-cloud that allows them to realise the full benefits, where different applications integrate seamlessly across clouds with data being transferred between them.

1.15 There are indications that competition for new customers is leading to some positive outcomes. Providers are investing in their offerings to match product development by their rivals and we see some evidence that they are responding to customer demand for open-source technologies, for example by adopting containers.[4] Customers also have access to a diverse range of services from ISVs, including some that meet very specialist use cases, that are developed and run using cloud infrastructure as the foundation.

## We are concerned about features of the market that create barriers to switching and multi-cloud

1.16 Given the complex nature of customer requirements in the cloud and technical variations between the solutions offered by different providers, there are always likely to be inherent barriers to switching and using multiple providers.

1.17 However, we have identified some features of the market that raise barriers to effective competition by making it more difficult for customers to switch and multi-cloud than might

[4] A container is a package of software that bundles an application's code with any necessary software required for the application to run (e.g. configuration files and libraries).

otherwise be the case. The features we are most concerned about are the charging of egress fees, technical barriers and the structure of committed spend discounts. We suspect that these practices, either alone or acting in combination, can limit the ability of customers to switch provider or adopt a more integrated multi-cloud strategy.

## The cost of egress fees can discourage customers from switching or using multiple cloud providers

1.18 Some cloud providers charge customers when they transfer data out of their cloud. This includes when they transfer data to end users and when they transfer data into a rival provider's cloud. These charges are known as egress fees. Egress fees can create significant additional cost and uncertainty for customers where they need to move data between providers on a regular basis. For example, where a business uses servers and storage in one cloud but wants to use the analytics service of a rival cloud that better suits its needs. Egress fees are also a commercial consideration when customers look to switch away from their existing cloud provider, particularly where they need to gradually move data and applications across to their new provider during the switching process.

1.19 Each of the hyperscalers charge a similar level of egress fees, which are around 5-10 times higher than some other cloud providers, such as OVHcloud and Oracle. Some cloud providers do not charge for egress at all. Our analysis indicates that egress fees at their current level are unlikely to be necessary for cost recovery and that egress list prices are likely to be higher than the incremental costs of providing the service.

1.20 Egress fees are a key concern for existing customers because they significantly increase the cost of taking a service from a different cloud provider. Our customer research found that 78% of respondents thought egress fees should be reduced or removed. We have heard examples where customers design their cloud architectures to intentionally avoid and reduce the cost of egress, which means they are unable to benefit from services from rival providers that may better suit their needs. This suggests that for some customers the costs associated with egress fees are likely to be significant enough to act as barrier to using multiple suppliers as part of a multi-cloud strategy.

## Technical barriers can limit the ability of customers to combine products from different providers or switch their main provider

1.21 The way different cloud services work together technically is a complex area that has a significant bearing on how competition works in cloud infrastructure. Where this works well, it can unlock significant benefits for customers by giving them access to the best products. However, a lack of interoperability and portability between services can result in customers needing to put additional effort into reconfiguring their data and applications so they can work on different clouds. This makes it more difficult to combine different services across cloud providers or to change primary provider.

1.22 Some of this complexity stems from technical differentiation between cloud providers, which can be the result of innovation which benefits customers. However, we are concerned that some of the barriers which arise from technical differentiation are not justified.

1.23 We have seen evidence of differences in the way AWS and Microsoft make the functionality of their cloud infrastructure services available when combined with their own services compared to those of competitors. Sometimes functionality is made available to competitor

services only after a delay, or in some cases not at all. Cloud providers, in particular AWS and Microsoft, may not always be fully transparent about the compatibility of their cloud infrastructure services with competing services from rivals, including ISVs.

1.24 Differentiation between providers for ancillary services (such as security, access management, monitoring and billing) may be greater than is necessary, thereby increasing complexity and cost of multi-cloud deployments. We have seen evidence that technical solutions (such as direct connection of data centres) exist to address the latency issues that can arise with multi cloud, but we find relatively little take up by the industry. While tools are available that facilitate switching and multi-cloud, we find these are limited and mostly focussed on hybrid cloud deployments - which combine on-premises and public cloud deployments - rather than between clouds.

1.25 Taken together, these barriers could limit the ability of customers to implement different multi-cloud architectures. This is likely to be most acute for customers with large numbers of applications or cloud architectures that are tightly integrated with many first-party proprietary services from their existing provider. These customers can find it more difficult to switch or build their preferred cloud architecture, where they can mix and match the cloud services that most closely meets their needs. Overall, we are concerned that technical barriers could dampen competition by lowering the threat of customers switching all or some of their workloads to benefit from better prices or higher quality cloud services.

## The structure of committed spend discounts can encourage some customers to use a single hyperscaler for most or all of their cloud needs

1.26 Committed spend discounts are when a customer agrees to spend a set amount with a single cloud provider in return for a percentage discount. They are usually part of an agreement between the leading providers and their larger customers. Customers with committed spend discounts account for a high proportion of the hyperscalers' UK revenues. An important feature of the discount structure is that the more a customer spends on the provider's cloud services, the greater the discount received.

1.27 Discounting can help customers to negotiate a good deal by committing to a set level of spend. However, the structure of these discounts acts as a barrier to muti-cloud by encouraging larger customers to use a single hyperscaler for all or most of their cloud needs. We have heard that this is an important commercial consideration for these customers, who feel discounting incentives encourage them to purchase most of their services from the same provider.

1.28 The prospect of receiving a lower discount can make it less attractive for customers to use a rival for some of their existing or new workloads.[5] We think this is a particular concern where customers face barriers to switching their existing cloud use. Ultimately this could restrict competition by raising barriers to entry and expansion for smaller cloud providers who cannot compete for customers with a broad set of cloud needs. It could also hamper the ability of rival providers to compete effectively for any new workloads as they emerge.

[5] A workload is a specific application, service, capability or a specific amount of work that can be run on a cloud resource.

## These barriers are likely to affect a material number of customers, especially those with more complex needs

1.29 Where there is active competition for new customers, in particular for larger businesses, those customers are likely to have a stronger bargaining position when first migrating to the cloud. However, after a customer makes the initial choice of cloud provider, in many cases AWS or Microsoft, they are more likely to deploy future workloads from within that ecosystem. We think this is partly explained by the barriers we have identified, which we consider are likely to be strong enough to result in a material number of customers having a limited ability to switch or use multiple providers.

1.30 The extent to which customers are affected by the barriers we have identified will depend on their individual needs. Some customers may be able to switch relatively easily as they take few products that are more easily ported between cloud environments (for example, basic IaaS products). Customers may also be able to reduce technical barriers to switching/multi-cloud to some extent by using container services or open-source services that are not specific to a particular cloud environment. In both cases, this is only likely to be feasible for the small number of customers with few applications and simple needs, such as smaller start-ups, and it comes with an additional cost.

1.31 Our evidence suggests that a large portion of the market has more complex needs and faces high barriers to switching or adopting more integrated multi-cloud architectures once they have chosen their primary provider. Large and more mature organisations are likely to be particularly affected. For example, these customers have large numbers of applications and/or use various proprietary services offered by their cloud providers, which add to the complexity of switching cloud provider.

## Limits on the ability to multi-cloud and switch can reduce competitive pressure on, and between, the market leaders

1.32 Where customers face material barriers to switching and multi-cloud, this can reduce competitive pressure on providers, as customers cannot credibly threaten to switch all or some of their existing workloads to a rival provider. We suspect that providers, and in particular AWS and Microsoft, hold a degree of market power in respect of the existing and incremental workloads of a material share of existing customers.

1.33 High levels of profitability for the market leaders AWS and Microsoft and a gradual increase in market concentration indicate there are limits to the overall level of competition. Our analysis indicates that AWS's profitability has been consistently high, with returns significantly above the WACC since at least 2014. We estimate that Microsoft's Azure returns have increased in recent years and are also above our estimate of the WACC. At the same time AWS's and Microsoft's share of the UK market has continued to increase, with their combined share of IaaS and PaaS revenues reaching 70-80% in 2022.

1.34 We are concerned that limits on competition create a significant risk of harm to cloud customers. This could lead to higher prices compared to what would be the case if customers could switch or multi-cloud more easily. Customers may also be harmed if there is a more innovative product on offer by a competitor and they cannot switch their existing

workloads. In principle, customers can protect themselves from these future risks when they initially contract with their provider, but in practice their ability to do so is limited due to factors such as the difficulty of forecasting their future demand for cloud services.

1.35 Harms for customers can translate into poor outcomes for UK consumers. Where businesses face higher costs of cloud infrastructure this will ultimately lead to higher prices for the products and services that they provide to consumers.

## We are concerned that the level of competition could deteriorate further over the longer-term

1.36 Looking ahead, we think there is a significant risk that the market becomes more concentrated as it matures, with less intense competition between the leading players.

1.37 Where customers have difficulty switching and using multiple providers, it could make it harder for smaller cloud providers to compete for those customers' workloads and grow their business as a result. In a maturing market where the number of new customers will reduce over time, this could make it more difficult for rivals to gain scale and challenge the market leaders effectively. This would be from a point where Microsoft and AWS have already established a strong position today. While it is difficult to predict what the exact market structure will look like in future, it is more certain that the outcome in this scenario would be further concentration around a small number of cloud providers.

1.38 Today we see some evidence that the market leaders have an incentive to compete to win new customers and to a much lesser extent for some narrow sets of additional workloads from existing customers. A weaker competitive constraint from rivals and barriers to switching and multi-cloud would allow the market leaders to entrench their position, while avoiding the need to compete intensely for each other's customers. This could reduce their incentive to discount prices or invest in developing services, either in response to competitive constraints from smaller providers or each other. With fewer new customers to compete for as the market matures, incentives to invest in innovation may reduce further.

1.39 In a more concentrated market, the leaders also have less incentives to support ISVs on their platform to attract new customers. We are concerned this could increase the ability and incentive of the market leaders to foreclose or exploit rival ISVs, for example by acting in ways that favour their own competing products. In turn, this impacts the choice, quality and prices that ISVs are able to offer to their customers.

## We are referring the cloud infrastructure market to the CMA for an in-depth investigation

1.40 Our study has found that, while there are some positive signs of competition at present, there are also clear indications that the cloud infrastructure market is not working well. We have identified features of the market that we think have an adverse effect on competition and could result in harm to customers and ultimately UK consumers. If left unchecked, we are concerned that these features could contribute to a further deterioration in competition in what is a critical market for digital services and the UK economy.

1.41 Ofcom may decide to refer a market to the CMA when we have reasonable grounds for suspecting that a feature or combination of features of a market or markets in the UK

prevents, restricts, or distorts competition. We consider that egress fees, restrictions on interoperability and committed spend discounts are barriers that make it more difficult for customers to change provider or use multiple suppliers. We have reasonable grounds to suspect that these features prevent, restrict or distort competition. We have also identified some credible interventions that could address the concerns we have identified.

1.42 On this basis we are referring the market for public cloud infrastructure services to the CMA to carry out a market investigation.[6] In reaching this decision we have assessed our concerns in line with CMA guidance on market investigations. Our assessment is that the legal threshold is met and a market investigation reference is an appropriate response to the concerns we have identified. We therefore exercise our discretion to do so. While we have identified some particular features of the market, it will be open to the CMA to investigate any other issues that it considers appropriate.

1.43 A market investigation reference is a significant step for us to take. Our decision reflects the importance of cloud computing to UK consumers and businesses and the significant concerns we have about the public cloud infrastructure market. The CMA will now conduct an independent investigation to decide whether there is an adverse effect on competition. Should it find an adverse effect on competition, the CMA will decide whether action should be taken to remedy, mitigate or prevent this or its detrimental effects on customers. The CMA has the ability to impose a broad range of remedies in response.

## We have also heard concerns about how software licensing practices could impact competition in cloud infrastructure

1.44 Some suppliers of cloud services have raised concerns with Ofcom regarding the software licensing practices of some cloud providers, particularly Microsoft. The concerns centre on the way Microsoft sells and licences some of its software products used by businesses. Among others, these include the Windows operating system, Microsoft SQL Server (a database management system) and the Microsoft 365 productivity suite (known as Office).

1.45 We have received submissions that say Microsoft engages in several practices that make it less attractive for customers to use Microsoft's licensed software products on the cloud infrastructure of rival providers compared to Microsoft Azure. The submissions allege that this limits their ability to compete for customers. Microsoft disputes the veracity of the concerns.

1.46 It is possible that the alleged conduct could risk dampening competition in cloud infrastructure services. We make no findings in relation to the complaints themselves in this report. It will be for the CMA to decide whether to investigate these issues further during the market investigation.

[6] Annex 6. Ofcom, 2023. Terms of reference.

# 2. Introduction

2.1 In this section, we provide some context to our market study including wider interest in cloud services competition, and digital markets more generally, at home and abroad. We also summarise the market study process, the evidence we have gathered over the course of our study and explain the purpose and structure of this report.

## Context

2.2 Cloud computing has become critical for many businesses across the economy – including telecoms companies, broadcasters and public sector organisations – and has transformed the way they deliver services on which we all rely every day. It uses data centres around the world to provide remote access to computing services such as software, storage and networking.

2.3 Demand for cloud services is growing and is expected to continue as the benefits become clearer and more widely accessible. We anticipate that dynamics in the markets for cloud services will be increasingly relevant for our duties in relation to competition, consumer protection, and network security and resilience in the communications sector. It is therefore important that we understand how these markets function and establish whether they are working well for consumers.[7]

2.4 On 6 October 2022, we launched a market study into cloud services, setting out our intention to gain a better understanding of this critical component of the digital economy, and to gather evidence to inform an assessment of whether competition is working well for consumers and citizens in the UK.[8]

2.5 In this study we have focussed on the market for public cloud infrastructure services which respondents broadly supported.[9] These services are the foundational elements of the cloud stack on which other cloud services (like software as a service, SaaS) are built.

2.6 On 5 April 2023, we published our interim report to provide an update on our approach and our progress with the study, to indicate the direction of travel our analysis was taking in relation both to concerns and potential interventions to address them, and to test these initial findings with stakeholders.[10] On the same day, we published a notice and consultation on a proposal to make a market investigation reference (MIR) to the Competition and Markets Authority (CMA) into the supply of public cloud infrastructure services in the UK.[11]

2.7 At the halfway stage of the study, we provisionally identified features and practices that make it more difficult for customers to switch and use multiple cloud providers. We were particularly concerned about the practices of Amazon Web Services (AWS) and Microsoft

[7] In September 2022, we published a document setting out our approach to competition and consumer issues in internet-based communications markets: Ofcom, 2022. Digital markets in the communications sector. Ofcom's approach to competition and consumer issues in internet-based communications markets.
[8] Ofcom, 2022. Cloud services market study notice and Cloud services market study. Call for inputs (CFI).
[9] Ofcom, 2022. Cloud services market study notice, paragraph 2 to 4 and CFI, paragraph 2.8 to 2.12.
[10] Ofcom, 2023. Cloud services market study. Interim report (interim report)
[11] Ofcom, 2023. Notice of a proposal to make a market investigation reference under section 131 of the Enterprise Act 2002. Cloud services and Public cloud infrastructure services. Consultation: Proposal to make a market investigation reference (MIR consultation).

because of their market position.[12] We proposed to refer the public cloud infrastructure services market to the CMA to carry out a market investigation. This would allow the CMA to further examine the nature and extent of barriers and if it finds an adverse effect on competition in relation to public cloud infrastructure services in the UK, consider whether action should be taken to remedy, mitigate or prevent the adverse effect on competition or its detrimental effects on customers.

2.8 This is the final report of our market study into the supply of cloud services in the UK.

## Wider UK policy and regulatory context

2.9 Throughout the study, we have engaged with other regulators, such as the CMA,[13] Information Commissioner's Office (ICO), [14] Prudential Regulation Authority (PRA), [15] and Financial Conduct Authority (FCA), [16] to feed into our understanding of the broader policy and regulatory landscape relating to cloud services and digital markets more generally in the UK. Ofcom is a member of the Digital Regulation Cooperation Forum (DRCF), which aims to facilitate coherence, cooperation, and collaboration between its members on digital regulatory matters.[17] During 2022/23, the DRCF supported effective and appropriate knowledge sharing on the wider regulatory landscape for cloud services.[18]

2.10 Relevant regulatory developments in the UK include the Bank of England, the FCA, and the PRA's consideration of the systemic risks that the reliance of UK financial institutions upon certain third parties (including cloud providers) raises to the stability or market confidence of the financial system of the UK.[19] The FCA is also examining the potential competition impacts of Big Tech entry and expansion in retail financial services.[20]

2.11 On 25 April 2023, the Government published the Digital Markets, Competition and Consumers Bill ("the Bill"). This Bill establishes a new pro-competition regime for digital markets and empowers the CMA to designate firms providing digital activities with strategic market status (SMS). Firms designated with SMS will be required to comply with conduct requirements to manage the effects of market power. The CMA may also apply pro-competitive interventions to tackle the root causes of an SMS firm's market power. The Bill contains measures to further support effective cooperation between the CMA and Ofcom in relation to communications matters. We considered how the new regime may be relevant to the firms considered in this market study and will continue to cooperate with the CMA,

[12] A combined revenue share in the UK of 60-70% in 2021.
[13] For further information about the CMA visit its website [accessed 10 August 2023].
[14] For more information on the ICO's responsibilities relating to cloud services, see: The Information Commissioner's response to Ofcom's cloud services market study call for inputs [accessed 4 August 2023].
[15] For further information about the PRA visit its website [accessed 10 August 2023].
[16] For further information about the FCA visit its website [accessed 10 August 2023].
[17] The DRCF brings together the CMA, FCA, ICO and Ofcom. For more information about the DRCF, see the Terms of Reference and its website [accessed 17 August 2023].
[18] See DRCF, 2023. Digital Regulation Cooperation Forum: Annual Report 2022/23 [accessed 4 August 2023].
[19] See FCA, 2023. Operational resilience: critical third parties to the UK financial sector [accessed 4 August 2023]. The Financial Services and Markets Act 2023 sets out the statutory framework for the supervisory authorities – FCA, PRA and Bank of England – to oversee the resilience of third parties services.
[20] See FCA, 2023. Feedback Statement: The potential competition impacts of Big Tech entry and expansion in retail financial services [accessed 4 August 2023].

along with the other relevant regulators, on the CMA's implementation of a new pro-competition regime for digital markets.[21]

2.12 The findings of this report will provide the Government and other UK regulators with a robust evidence base and a broad understanding of the market for cloud services and, in particular, public cloud infrastructure services, supporting collaboration on regulatory approaches in the UK.

# International context

2.13 Cloud services are of interest to several jurisdictions around the world given their role in the global digital economy. A summary of relevant regulatory developments outside the UK is provided below.

## European Union (EU)

2.14 The cloud sector is affected by both existing and emerging legislative developments within the EU. Notable examples include:

a) **Digital Markets Act (DMA)**: Largely applicable from 2 May 2023, this imposes a suite of ex ante regulatory obligations on large digital platforms that meet the requirement of a "gatekeeper" for one or more "core platform services", which includes "cloud computing services." Designated gatekeepers are subject to a mixture of obligations such as a duty to ensure data portability to their users and a prohibition of self-preferencing.[22] However, no digital platform has so far been designated a gatekeeper with respect to cloud computing services in the first tranche of designations.[23]

b) **Data Act**: Political agreement on the Data Act was reached on 27 June 2023. The legislation (in the form of a Regulation) seeks to set out the rules on who can use and access what data (and on what terms) generated across all economic sectors in the EU. The proposed rules include allowing customers to switch effectively between different "data processing service providers" (including cloud service providers). Most notably the Regulation contains a provision for the abolition of cloud switching charges (including data egress charges) after a transition period of three years after the date of entry into force of the Regulation, as well as provisions aiming to improve interoperability for in-parallel use of multiple data processing services.

## France

2.15 On 29 June 2023, the French competition authority, the *Autorité de la concurrence,* published its market study into competition in the cloud sector.[24] *The Autorité de la concurrence* analysed a number of practices implemented, or likely to be implemented in the sector, that it felt could restrict competition on the merits e.g. cloud credits and egress

[21] See DRCF, 2023. Digital Regulation Cooperation Forum: Workplan 2023/24 [accessed 4 August 2023].
[22] Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act) [accessed 27 September 2023].
[23] European Commission, 2023. Press Release. Digital Markets Act: Commission designates six gatekeepers [accessed 12 September 2023].
[24] *Autorité de la concurrence*, 2023. Press release. Cloud computing: the Autorité de la concurrence issues its market study on competition in the cloud sector | Autorité de la concurrence [accessed 25 August 2023].

fees. It concluded that competition in the cloud industry was characterised by "competition *for* the market rather than *on* the market, insofar as, for a specific need or workload, customers tend to turn to a single supplier, particularly those with an attractive ecosystem".[25] The *Autorité de la concurrence* has stated that the competitive risks outlined in its opinion will be analysed by its investigation teams and notes the different regulatory and competition law tools at its disposal to tackle any restrictive competition practices.

## The Netherlands

2.16 On 5 September 2022, the Dutch competition authority, Authority for Consumers and Markets (ACM), published its market study into cloud services.[26] The ACM found that it was difficult for smaller players in the Dutch market to compete effectively with large integrated providers.[27] This was perpetuated by "vendor lock-in",[28] in part reinforced by "poor interoperability"[29] and other barriers to switching.

2.17 Following the completion of its market study, the ACM launched a follow-up investigation into competition problems caused by barriers to switching cloud providers. However, the ACM elected to close this investigation as it felt that both the Data Act and the DMA will be able to "solve several major problems sooner". The ACM said that it reserves the right to launch another investigation under its competition rules should new evidence come to light with respect to switching barriers erected by providers.[30]

## Japan

2.18 On 28 June 2022, the Japanese competition authority, the Japan Fair Trade Commission (JFTC), published the findings of its fact-finding survey regarding both trade practices and the state of competition in the cloud services sector.[31] The report highlighted the type of conduct that might restrain competition in the market and made some recommendations to both suppliers and customers of cloud services with respect to actions that they could respectively take to encourage a competitive market.

## United States (US)

2.19 The US Federal Trade Commission (FTC) issued a request for information (RFI) on 22 March 2023 seeking information about the competitive dynamics of cloud computing, the extent to which certain segments of the economy are reliant on cloud service providers, the security risks associated with the industry's business practices, and the interactions between AI and cloud computing.[32] The RFI closed for comments on 21 June 2023, and the FTC is currently analysing the responses[33] received.

---

[25] *Autorité de la concurrence*, 2023. Summary of Opinion 23-A-08 of 28 June 2023 on competition in the cloud sector, page 13 [accessed 25 August 2023].
[26] ACM, 2022. Market Study Cloud services [accessed 25 August 2023].
[27] Ibid, page 6.
[28] Ibid, page 5.
[29] Ibid, page 5.
[30] ACM, 2023. Press Release. European Data Act to make it easier to switch cloud services [accessed 25 August 2023].
[31] Japan Fair Trade Commission, 2022. Report Regarding Cloud Services [accessed 25 August 2023].
[32] FTC, 2023. FTC Seeks Comment on Business Practices of Cloud Computing Providers that Could Impact Competition and Data Security [accessed 25 August 2023].
[33] https://www.regulations.gov/docket/FTC-2023-0028/comments [accessed 20 September 2023].

# The market study process

2.20 Ofcom has concurrent functions with the CMA pursuant to section 370 of the Communications Act 2003. This includes the power to undertake a market study to consider the extent to which a matter in relation to commercial activities connected with communications matters has or may have effects adverse to the interests of consumers.[34]

2.21 Market studies are examinations into the causes of why particular markets may not be working well and in the interests of consumers, taking into account any regulatory and economic drivers, and patterns of suppliers' and customers' behaviour.[35]

2.22 Further information on market studies can be found in the following guidance documents: Market Studies: Guidance on the OFT Approach (OFT519)[36] and Market Studies and Market Investigations: Supplemental Guidance on the CMA's Approach (CMA3).[37]

2.23 Key milestones for the cloud service market study:

a) We published a market study notice on 6 October 2022 launching a study into the provision of cloud services in the UK.[38]
b) We published a notice on 5 April 2023 together with a consultation document inviting views on our proposal to make a MIR into the supply of public cloud infrastructure services in the UK.[39]
c) This report, published on 5 October 2023, sets out our findings and our decision to make a MIR to the CMA together with our reasons and supporting information. The terms of reference for the MIR are included in this report at Annex 6. We have also published it alongside this report on our website.[40]

2.24 We consulted with the CMA as we developed plans to carry out a market study on cloud services and throughout the course of this study.[41]

# Evidence gathering

2.25 Over the course of the study, we have gathered information from different sources as summarised below.

## Call for inputs (CFI)

2.26 We began our study by publishing a CFI seeking views on:

a) the proposed scope of the study;

[34] Communications matters includes services made available by means of or to facilitate the provision of electronic communications networks and/or electronic communications services. See section 369(1) of the Communications Act 2003.
[35] "Customers" include businesses as well as residential consumers. See section 183(1) the Enterprise Act 2002.
[36] Office of Fair Trading, 2010. Market studies: Guidance on the OFT approach [accessed 11 August 2023].
[37] CMA, 2017. Market Studies and Market Investigations: Supplemental guidance on the CMA's approach [accessed 11 August 2023].
[38] Ofcom, 2022. Cloud services market study notice.
[39] Ofcom, 2023. Notice of a proposal to make a market investigation reference under section 131 of the Enterprise Act 2002. Cloud services and Public cloud infrastructure services. Consultation: Proposal to make a market investigation reference.
[40] Ofcom, 2023. Terms of reference.
[41] As provided for by section 370(5) of the Communications Act 2003.

b) our initial characterisation of the market;
c) our proposed approach for considering the dynamics in cloud infrastructure services competition and to examine cloud ecosystem competition;
d) any concerns regarding any conduct or activities of any provider(s) that may adversely affect market dynamics now or in the future; and
e) any remedies that we should investigate further to mitigate some of the potential risks or concerns with the market.

2.27 We received 11 responses to our CFI and published non-confidential responses on our website.[42]

## Market research

2.28 To help us better understand the customer perspective, we commissioned some market research. We published a summary of findings and accompanying data tables on 5 April 2023 alongside our interim report (see 'Interim report and MIR consultation' below).[43] This research included both qualitative and quantitative phases. The research involved 50 one-hour discussions and over 1000 survey interviews with UK decision-makers in UK businesses that used, or were considering using, IaaS and/or PaaS services. The research included a range of company size bands and industry sectors.

## Customer engagement

2.29 We supplemented our market research by also gathering information from 12 of the UK's biggest cloud customers who responded to a questionnaire which we sent to them. We also met many customers from different sectors of the economy (private and public) to hear directly from them about their experiences in choosing, buying and using cloud services. We also spoke with trade associations and academics with an interest in cloud services.

## Telecoms and broadcasting

2.30 Given our sectoral interests, we have had discussions with some telecoms providers and broadcasters about their current and future use of cloud services and used our statutory powers to request information from them.[44]

## Suppliers in the cloud services value chain

2.31 Turning to the supply-side, we have had discussions with the hyperscalers[45] and other cloud providers. We have used our statutory powers to gather information from them. We have also spoken to and gathered information from other players in the value chain such as independent software vendors (ISVs) and providers of professional services such as resellers, consultants, and managed service providers.

[42] Non-confidential responses to our CFI are published on our website.
[43] Context Consulting, 2023. Cloud Services Market Research, Summary of Findings and Data Tables.
[44] Section 174 of the Enterprise Act 2002.
[45] Amazon Web Services (AWS), Microsoft and Google.

## Market analysts and public information

2.32 We purchased insights from some market analysts and reviewed publicly available information and literature about cloud services.

## Interim report and MIR consultation

2.33 Within six-months of launching our study, we published an interim report and proposed MIR consultation.

2.34 We received responses from 26 stakeholders to these consultations. Non-confidential responses are published on our website.[46]

## Webinar

2.35 On 10 May 2023, we presented our interim findings by means of a webinar to over 150 external stakeholders. We did this to raise awareness of, and stimulate wider engagement and input into, our study particularly among ISVs and cloud customers.

# Our final report

2.36 This final report sets out, among other things, our findings in relation to the matters specified in our market study notice. It also sets out our proposals on how to address any concerns we have identified and the next steps beyond the study as provided for by legislation. This report is structured as follows:

a) **Section 3** provides context about the cloud services market, providing an overview of the service and deployment models, and outlines the different players in the market.
b) **Section 4** details how competition works in the sector. It considers the customer perspective, how providers compete and key market outcomes (such as market shares and profitability).
c) **Section 5** considers the extent of barriers faced by customers who wish to use multiple clouds, multiple vendors and switch between them.
d) **Section 6** considers the extent of barriers to market entry and expansion.
e) **Section 7** considers the relationship between hyperscalers and ISVs.
f) **Section 8** summarises our findings on the current state of competition.
g) **Section 9** summarises submissions we have received regarding Microsoft's software licensing practices and articulates the relevance of those submissions for competition in cloud infrastructure.
h) **Section 10** identifies potential intervention options which could address our concerns about the market.
i) **Section 11** sets out our decision to refer the market for public cloud infrastructure services to the CMA for further investigation.

2.37 The following supporting annexes also form an integral part of the report:

a) **Annex 1** sets out our analysis of revenues and shares of supply associated with cloud infrastructure services in the UK, as well as the types of products which generate most revenue for hyperscalers in the UK.

[46] Responses to our interim report and responses to our MIR consultation.

b) **Annex 2** sets out our analysis of the profitability of the hyperscalers' cloud infrastructure services. We compare hyperscaler operating profits and margins to those of other cloud providers, and we assess whether the market leaders' cloud businesses have generated returns persistently above their cost of capital.
c) **Annex 3** sets out the evidence we have reviewed on the prevalence of multi-cloud and switching, including our assessment of the hyperscalers' submissions on this.
d) **Annex 4** sets out further detail on the technical barriers we have identified in Section 5.
e) **Annex 5** sets out our assessment of the evidence received by the hyperscalers on discount outcomes, which mainly focused on cross-service privately negotiated committed spend customers, as well as the evidence received from customers on their experience of negotiating with the hyperscalers.
f) **Annex 6** sets out the terms of reference for the market investigation.
g) **Annex 7** provides a glossary of terms used in the report.

# 3. Market context

## Introduction

3.1 In this section we set out the market context for our market study. We explain what cloud computing is and summarise the main cloud services, service models and deployment models. We highlight the importance of cloud services in the UK and provide some findings from our market research with cloud customers. We also introduce the key players in the cloud market and outline the role of cloud services in the telecoms, broadcasting and public sectors.

## What is cloud computing?

3.2 In this market study, we define cloud computing as the provision of remote access to computing resources (compute, storage and networking) on demand and over a network (public internet or a private connection), instead of a personal computer or local server that are not part of the cloud.

3.3 The National Institute of Standards and Technology (NIST)[47] in the US defines cloud computing as "a model for enabling ubiquitous, convenient, and on-demand network access to a shared pool of configurable computing resources (e.g. networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."[48] The UK Government offers a similar definition of cloud services: "a digital service that enables access to a scalable and elastic pool of shareable computing resources".[49]

3.4 The above definitions focus on the public cloud deployment model, where cloud services offer access to a shared pool of computing resources. However, alongside public cloud, there are two additional deployment models: private cloud, where the computing resources are not shared between customers, and hybrid cloud, which combines aspects of public and private cloud. Regardless of the delivery model, cloud computing is distinct from traditional IT where assets are usually located on site and are not part of the cloud.

3.5 Traditional IT infrastructure is made up of data centres, servers, networking hardware, desktop computers and applications. It is usually installed on-premises for private use by an organisation. It is usually connected to a network which includes stored data and applications. Organisations relying on traditional IT infrastructure normally depend on an in-house IT department to install and maintain the infrastructure.

3.6 Compared to traditional IT infrastructure, cloud computing offers flexibility and scalability which enables customers to quickly scale up or down the computing resources that support their business. This can allow them to reduce their IT costs, transform capex into opex,

[47] The National Institute of Standards and Technology is a non-regulatory agency within the US Department of Commerce. NIST website [accessed 12 September 2023]. The NIST definition of cloud computing is widely adopted.
[48] NIST, 2011. The NIST Definition of Cloud Computing [accessed 28 July 2023].
[49] DCMS, 2022. Policy paper. Data storage and processing infrastructure security and resilience - call for views [accessed 28 July 2023].

increase their innovation potential, enhance their quality of service, and achieve baseline security and resilience.[50] It also offers access to relevant data from any device, anytime and anywhere.

3.7 While cloud computing offers significant benefits, it does have some limitations. Cloud infrastructure is usually owned and managed by the cloud provider, so the customer may have more limited control over their data, applications and services. Furthermore, public cloud computing is completely reliant on internet connection, so if the connection is interrupted, data cannot be accessed.

3.8 The main suppliers of cloud services (all services involved in the provision of cloud computing) in the UK are Amazon Web Services (AWS), Microsoft and Google,[51] which offer a broad range of cloud services at scale and are often referred to (in this document and more widely) as the hyperscalers. There are a number of smaller suppliers of cloud services, some offering a broad range of cloud services, while others are more specialised. We discuss suppliers in more detail later in this section.

## Importance of cloud services in the UK

3.9 Cloud services are increasingly important inputs to many businesses and organisations across the economy. Cloud computing supports not only the communications sector, but most other sectors, for example manufacturing, retail, hospitality and financial services, plus public and voluntary sector bodies. Without cloud many digital businesses providing services to consumers would not be able to function in the way they do today. A well-functioning cloud market is essential to UK productivity today and in the future.

3.10 Cloud is also a cornerstone of recent technological innovations. From data science to AI, many of the cutting-edge developments in the way software is transforming how we live our lives, run our businesses, and engage with our public services, rely on the cloud. In particular, we expect that cloud computing will underpin the development of AI as it provides the computing resources and infrastructure needed to train and deploy AI models at scale. Related to this, the CMA has recently published a report on AI foundational models. It explores foundation models and how their use could evolve; what opportunities and risks these could bring; and some proposed principles that can help guide the development of these markets going forward.[52]

## Market maturity

3.11 Worldwide end user spending on public cloud services is already valued at £397 billion in 2022 and forecast to grow 22% to total £483 billion in 2023.[53] In comparison, worldwide IT

[50] AWS claims that customers moving to the cloud can achieve on average 80% reduction in IT carbon emissions. Amazon's public messaging on sustainability, including more detail on the 80% figure, is available at: AWS website. Sustainability, Innovating Products and Services, The Cloud [accessed 28 July 2023]. Capex or capital expenditures are major purchases a company makes that are intended to be used over the long term. Opex or operating expenses are the routine expenses a company incurs to remain operational.
[51] We use AWS, Microsoft and Google as they are the direct providers of the three hyperscaler clouds in the UK: AWS, Azure and Google Cloud. AWS is a subsidiary of Amazon and Google is a subsidiary of Alphabet.
[52] CMA, 2023. AI Foundational Models: Initial Report [accessed 19 September 2023].
[53] Gartner, 2023. Press Release: Gartner Forecasts Worldwide Public Cloud End-User Spending to Reach Nearly $600 Billion in 2023. [accessed 25 August 2023]. These figures have been converted into pound sterling using ONS exchange rate data (ONS Average Sterling exchange rate: US dollar).

spending is projected to total £3.8 trillion in 2023, an increase of 4% from 2022.[54] For the UK, International Data Corporation (IDC) projects that spending on public cloud services will similarly grow by 22% in 2023, while total IT spending will grow by 5% in 2023.[55] Based on the available data, public cloud spending is projected to continue on a growth trajectory.

3.12 Many businesses and organisations are at some stage of modernising their IT through the adoption of cloud computing. Evidence suggests that many customers have migrated large parts of their workloads to the cloud, although levels of adoption are likely to vary by sector.[56] For example:

a) Our customer research found that 82% of respondents had increased their spend on cloud in recent years, with 26% having greatly increased their budget.[57]
b) Flexera research in 2023 with (mainly larger) European organisations found that more than half (62%) of respondents stated they were using cloud heavily.[58]
c) Gartner 'cloud shift' research in 2022 forecast that enterprise IT spending on public cloud computing, within addressable market segments, will overtake spending on traditional IT in 2025.[59]

3.13 We expect cloud services to become even more important in the next few years. Some late adopters may begin the migration to cloud, while research suggests existing cloud users will transition more of their workloads from on-premises as new use cases emerge:

a) Our customer research found that 79% of respondents expect to spend more on cloud in the next 18 months.[60]
b) A 2022 paper by [✂] noted that a survey of organisations from the UK and US found respondents estimated that 40% of their workloads are in the public cloud today and estimated that 70% of workloads would be in public cloud in the next three years.[61]
c) A Goldman Sachs survey of IT executives from 100 Global 2000 companies in June 2022 suggested that 24% of respondents' workloads were already on the public cloud, with 42% of their workloads expected to be on the public cloud in the next three years.[62]

---

[54] Gartner, 2023. Gartner forecasts worldwide IT spending to grow 4 percent in 2023 [accessed 25 August 2023]. These figures have been converted into pound sterling using ONS exchange rate data (ONS Average Sterling exchange rate: US dollar).
[55] IDC, 2023. Worldwide Black Book: Live Edition, July (V2 2023) Forecast (published July 2023). Total IT spending includes the following IDC technology categories from the Black Book publication: Infrastructure, Application Development & Deployment, Applications, System Infrastructure Software, Managed Services, Support Services, Project Oriented Services and Devices.
[56] Hewlett Packard Enterprise analysis of responses to its freedom of information request to UK public sector technology professionals found that as much as 70% of public sector organisations' infrastructure and 73% of data remains on premises. Hewlett Packard Enterprise, 2022. Public Sector Cloud Strategy Report [accessed 12 September 2023].
[57] Context Consulting research report, slide 39. The sample covered decision-makers for UK organisations that were existing users of cloud computing services (IaaS, PaaS or both) or those considering adoption within 12 months.
[58] Flexera, 2023. 2023 State of the Cloud Report, page 75 [accessed 29 September 2023]. It is worth noting that a large share (62%) of these European respondents were from the UK.
[59] Gartner, 2022. Gartner Says More Than Half of Enterprise IT Spending in Key Market Segments Will Shift to the Cloud by 2025 [accessed 28 July 2023].
[60] Context Consulting research report, slide 42.
[61] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[62] Goldman Sachs Equity Research, 12 July 2022. IT Spending Survey.

d) IDC estimates that organisations' spend on cloud services accounted for around 30% of total IT spending in 2022 – both in the UK and globally.[63] IDC projects that the cloud deployment share of total IT spending will rise to 47% globally by 2027, with a slightly faster rise to 51% in the UK by 2027.[64]

3.14 Based on all of the above and evidence gathered during the course of our market study, it is clear that that UK business and organisations have already started migrating workloads to cloud to differing degrees, and we expect that trend to continue. Ideally, we would want to understand how the remaining IT workloads are likely to migrate to the cloud (noting this will not be all of them) and the split between i) existing customers migrating new or additional workloads, and ii) completely new cloud customers. However, we have not seen such a statistic. Based on our understanding, we think that in the future there will be relatively fewer completely new customers, and that most new workloads (which may still be a significant number, as per the projections above) will be from existing customers.

# What are cloud services?

3.15 Cloud services provide access to computing resources on demand, via a network. The customer buys access to the computing resources as a service and typically does not own the underlying hardware and software. There are three key elements to this definition:

a) Computing resources – these include hardware (servers and network equipment) and software (applications) which are used to process workloads[65] and store data.
b) On demand – the computing resources are available on a scalable and elastic basis. This typically involves the dynamic provision of virtualised computing resources. Users are often billed for the amount of resource used.
c) Via a network – the transit of data to and from the cloud provider may be over the public internet or a private connection. This allows location-independent access to the cloud.

3.16 Cloud services started to be used at scale when they were launched by AWS, followed by Microsoft and Google. Originally, AWS cloud services were used internally to support Amazon's online retail services. In 2006, AWS officially launched its cloud services for third party use.[66] Some years later Microsoft followed suit with Microsoft Azure in 2010 and Google with Google Cloud in 2011.[67]

## Service models

3.17 Cloud services are typically classified according to their service models: infrastructure as a service (IaaS), platform as a service (PaaS) and software as a service (SaaS). These three

[63] IDC, 2023. Worldwide Black Book: Live Edition, July (V2 2023) Forecast (published July 2023). Total IT spending includes the following IDC technology categories from the Black Book publication: Infrastructure, Application Development & Deployment, Applications, System Infrastructure Software, Managed Services, Support Services, Project Oriented Services and Devices.
[64] IDC, 2023. Worldwide Black Book: Live Edition, July (V2 2023) Forecast (published July 2023).
[65] A workload is a specific application, service, capability or a specific amount of work that can be run on a cloud resource.
[66] AWS website. About AWS [accessed 28 July 2023].
[67] AWS, Microsoft and Google offered some cloud services in beta version before the official launch of their clouds.

service models have been recognised by the NIST.[68] It is worth noting that some services may not 'fit' neatly into these service models and their most suitable 'placement' is the topic of ongoing discussion in the wider professional community, though we still consider them to be useful to inform our analysis in this market study.

3.18 The service models are differentiated by the level of control the customer has over the management and maintenance of the computing resources.[69] IaaS, PaaS and SaaS form a vertical stack, where each layer is notionally built on top of the previous one(s). This is shown in Figure 3.1 below.

**Figure 3.1: The cloud computing stack**

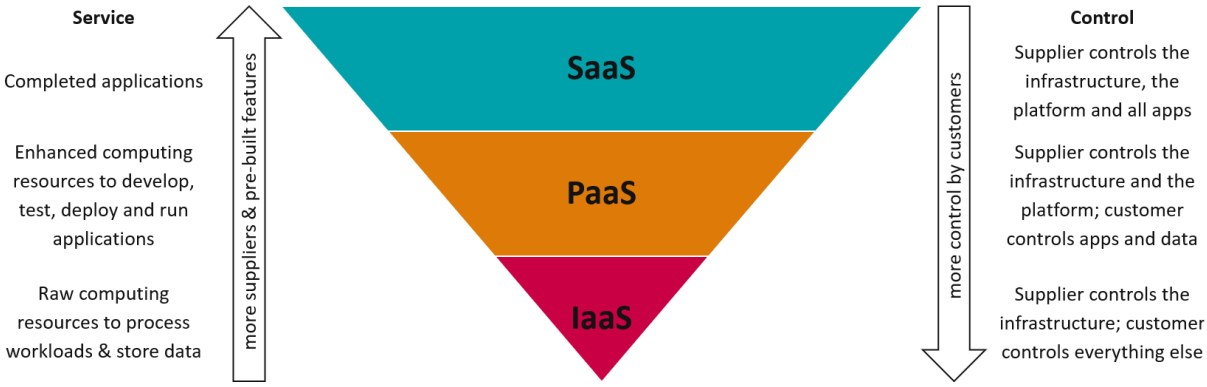


*Source: Ofcom.*

3.19 **Infrastructure as a service (IaaS)** are cloud services that provide access to raw computing resources for processing workloads and storing data. These computing resources are in the form of servers and networking equipment owned and managed by the IaaS provider (and typically held on racks in a remote data centre). To allow and manage that access, IaaS also includes some necessary software, including networking (e.g. firewall) and virtualisation.[70] The customer has the highest level of control over the cloud stack, including over the operating system, applications and data. Examples of IaaS include AWS EC2, Microsoft Azure Virtual Machines and Google Compute Engine – which can be used by business customers, for example, to store data and install software.[71] IaaS should be distinguished from **bare metal** services, which offer access to dedicated servers with no or limited software installed (e.g. no operating system or virtualisation). We estimate that in 2022, UK IaaS revenues were around £[✂] [£4.0 to £4.5] billion and grew by 30% to 35% per year between 2019 and 2022.[72]

[68] NIST, 2011. The NIST Definition of Cloud Computing [accessed 28 July 2023].
[69] Control refers to the involvement the customer has in the management and maintenance of the computing resources themselves, as opposed to the freedom it affords them to, for example, choose between providers. We will assess this separately as discussed elsewhere in this document.
[70] Virtualisation is the process of using software to create an abstraction layer over servers that allows the hardware elements of a single server (e.g. central processing units, random access memory and storage) to be divided into multiple virtual servers, commonly called virtual machines.
[71] AWS website. Amazon EC2 [accessed 28 July 2023]; Microsoft Azure website. Virtual Machines [accessed 28 July 2023]; and Google Cloud website. Compute Engine [accessed 28 July 2023].
[72] Ofcom analysis of data provided in response to our information requests and data from Synergy and IDC. Annual growth based on the compound annual growth rate between 2019 and 2022. See Section 4 on UK shares of supply for more detail.

3.20 **Platform as a service (PaaS)** are cloud services that provide access to a virtual environment for customers to develop, test, deploy and run applications. These include application development computing platforms and pre-built application components and tools which customers can then use to build and manage full applications. There are many PaaS products, and key categories include databases, analytics, containers,[73] machine learning and IoT (internet of things). The overall virtual environment and the underlying raw computing resources are typically owned and managed by the same cloud provider.[74] However, the individual PaaS services (computing platforms, and/or pre-built application components and tools) may be supplied by the cloud provider or by independent software vendors (ISVs). The customer has less control over the cloud stack compared to IaaS; they still manage applications and data, but not the PaaS computing platform (including its operating system) or the pre-built application components and tools. Examples of PaaS products include AWS Elastic Beanstalk, Microsoft Azure DevOps and Google App Engine – which can be used, for example, to build streaming video on demand (SVoD) services.[75] We estimate that in 2022, UK PaaS revenues were around £[✂] [£2.5 to £3.0] billion, and grew by 35% to 40% per year between 2019 and 2022.[76]

3.21 **Software as a service (SaaS)** are complete applications hosted in the cloud. These cloud applications can be offered by the cloud provider that owns the underlying raw computing resources or by an ISV. The provider of the SaaS service manages all hardware and software. In general, most modern consumer and business facing applications are SaaS, including communications services (e.g. Gmail and WhatsApp), broadcasting video on demand (BVoD) services (e.g. BBC iPlayer), productivity software (e.g. Microsoft Office 365 and Google Workspace) and customer relationship management software (e.g. Salesforce Sales Cloud). Estimates of the size of the UK market for SaaS vary given difficulties determining the boundaries of SaaS, but it is likely to be larger than public IaaS and public PaaS combined.

[73] A virtual machine is a software-defined computer that is created by running a guest operating system on top of the host operating system of the physical server. Each virtual machine runs its own operating system and behaves like an independent server, even though it is running on just a portion of the actual underlying server hardware. The software that creates, runs and manages virtual machines is called a hypervisor. A container is a package of software that bundles an application's code with any necessary software required for the application to run (e.g. configuration files and libraries). Virtual machines and containers offer similar functionalities, but containers are typically lighter because they do not need to run a full operating system.
[74] There are examples of PaaS where the service provider owns the virtual environment but not the underlying raw computing resources. For example, IBM Red Hat OpenShift and VMware Tanzu are PaaS virtual environments that can integrate with many clouds, including those of the hyperscalers.
[75] AWS website. AWS Elastic Beanstalk [accessed 28 July 2023]; Microsoft Azure website. DevOps solutions on Azure [accessed 28 July 2023]; and Google Cloud website. App Engine [accessed 28 July 2023].
[76] Ofcom analysis of data provided in response to our information requests and data from Synergy and IDC. Annual growth based on the compound annual growth rate between 2019 and 2022. See Section 4 on UK shares of supply for more detail.

**Figure 3.2: Vertical stack for traditional IT and cloud computing**

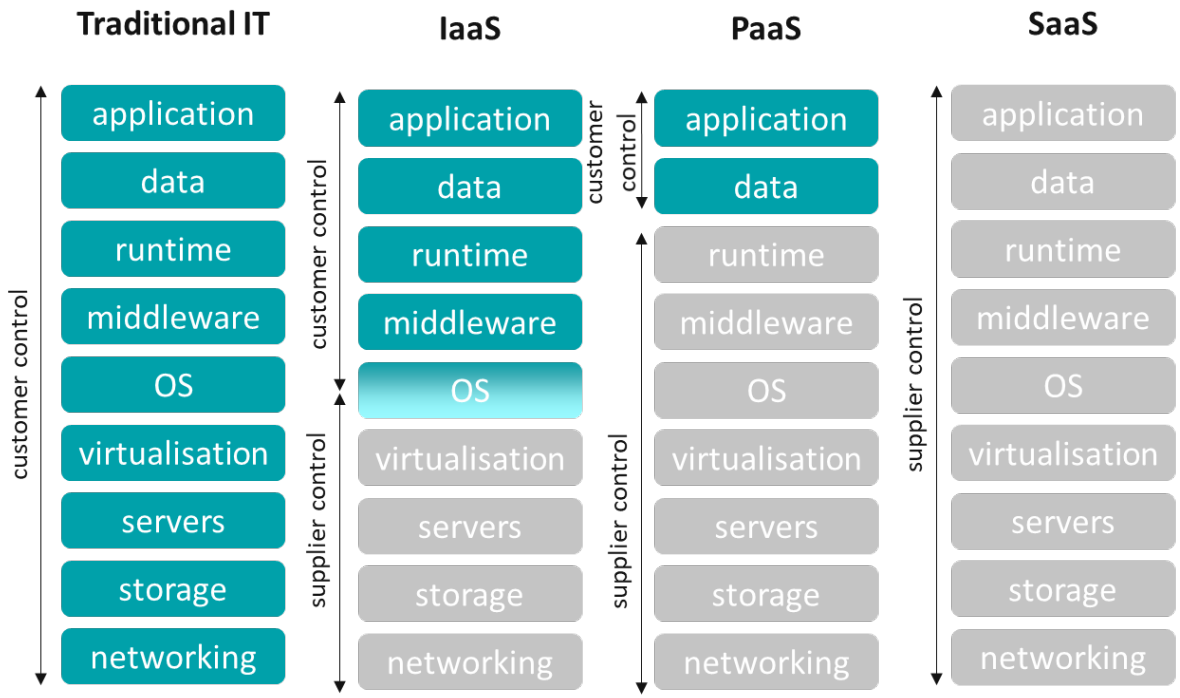


*Source: Ofcom.*

3.22 Some suppliers[77] of cloud services do not always use the above-described service models in their commercial offerings to customers, and instead prefer to group their services by the type of computing capability that they offer. This may reflect the fact that, within layers, there can be varying applications with different levels of control. Potentially, anything can be offered as a service, leading to the designation 'anything as a service.' This way cloud services can be split into many categories, including virtual machines, storage as a service, container as a service (CaaS), database as a service (DBaaS) and disaster recovery as a service (DRaaS) – all of which comprise a combination of cloud services from the three service models set out above.[78] Nevertheless as above, we still consider them to be useful to inform our analysis in this market study.

## Deployment models

3.23 Cloud deployment models indicate how the cloud services are made available to customers. Usually, they are classified into three major groups, namely:

a) Public cloud is the most common cloud deployment model, where cloud services are open to all customers willing to pay and computing resources are shared between them. Public cloud servers are typically located in an off-premises data centre and accessed remotely over the public internet or via dedicated connections. Customers of public

[77] e.g., AWS said in its response to the interim report that customers are typically looking to solve a specific IT problem and do not tend to make purchasing decisions based on the categorisation of a specific service as a IaaS or PaaS (AWS response to the interim report, page 2).

[78] Microsoft Azure website. Virtual Machines [accessed 28 July 2023]; IBM website. IBM Storage as-a-Service [accessed 28 July 2022]; Red Hat website. What is CaaS? [accessed 28 July 2023]; Oracle Cloud Infrastructure (OCI) website. What is DBaaS? [accessed 28 July 2023]; and VMware website. What is disaster recovery as a service (DRaaS)? [accessed 28 July 2023].

cloud services are typically businesses whose demands vary over time and buy cloud services on a pay-as-you-go (PAYG) basis.

b) Private cloud is a cloud deployment model where computing resources are dedicated to (as opposed to shared between) individual customers. It combines many of the benefits of cloud computing with the security and control of traditional IT. Customers may choose to use private cloud for various reasons, including in cases where their traditional IT is not easily transferable to the public cloud and for running latency-sensitive workloads close to them. Private cloud comes in many forms: it could involve the exclusive allocation of physical or virtual computing resources, it could be deployed in remote data centres or on the premises of the customer, and it could be provided by a third party or self-supplied.

c) Hybrid cloud is a cloud deployment model involving a combination of public clouds and private environment, such as private clouds or on-premises resources, which allow workloads to be shared between them.

## Multi-cloud

3.24 Customers sometimes purchase cloud services from more than one supplier which is often referred to as multi-cloud. In the course of our study, we have not come across a standard definition of the term multi-cloud and have found it can be used to mean different things. Accordingly, for the purpose of our study given the range of different ways in which customers purchase services from different suppliers we have adopted the following terms:

a) 'Multi-cloud': the take-up cloud infrastructure services offered by different cloud providers.

b) 'Multi-vendor': when first-party and ISVs services are used within the same cloud.

3.25 The use of multiple public clouds can benefit customers by allowing them to access their preferred services, gain commercial bargaining power against their cloud providers and build for resilience.

3.26 In practice, multi-cloud may take various forms and for our purposes can be broadly categorised as follows:[79]

a) **Cloud duplication**: this occurs whenever customers aim to mirror their cloud architecture on two or more public clouds, so that all or some of their applications and data can run equivalently on all of them. This architecture appears to be relatively rare and is mostly implemented for resilience by duplicating some parts of customers' cloud architecture (as opposed to the full architecture) to maximise service availability for critical applications in case of outage. That said, fully duplicating functionalities across several public clouds is common for ISVs offering cloud infrastructure services, as they may seek to deploy their services on different public clouds to access a larger pool of potential customers.

b) **Siloed multi-cloud**: this occurs where the customer runs different customer applications, stores different customer data sets and/or uses different cloud services hosted on two or more public clouds with no or minimal integration between these clouds (i.e.

[79] The definition of multi-cloud presented here slightly differs from the one set out in our CFI (at paragraph 3.35). We note there are various ways to define and/or categorise multi-cloud and there is not one standard definition. However, we consider the categorisation presented here matches with our understanding of multi-cloud based on the evidence we have gathered and assessed, and it is useful for our purpose of assessing the state of competition in public cloud infrastructure services.

different applications are 'siloed' on different public clouds). This multi-cloud architecture appears to be the most common amongst customers who use more than one cloud provider. It may be implemented to access 'best-of-breed' functionalities across clouds for specific applications but may also reflect an independent process of cloud uptake by different units of the same company (e.g. different departments may have independently migrated to different cloud providers).

c) **Integrated multi-cloud**: this multi-cloud architecture occurs where customers build their cloud architectures by mixing and matching cloud services hosted on different public clouds. [80] There is a spectrum of how integrated such architectures can be – anywhere from partially integrated (e.g. a second cloud is used only for a specific service or application) which is closer to siloed multi-cloud, to highly integrated (e.g. integrating multiple applications and data hosted on different clouds), to extremely integrated (e.g. dynamic distribution of microservices across different clouds). As an example, a customer may wish to run a data analytics application by integrating services hosted on Azure (e.g. compute and storage) and Google (e.g. BigQuery). This multi-cloud architecture may be implemented to access 'best-of-suite' functionalities across public clouds. However, as discussed in later sections, it can present customers with high financial and technical costs which may explain its limited uptake.

3.27 We consider the prevalence of multi-cloud in more detail in Section 4 and Annex 3, and we look at the potential technical barriers to multi-cloud in Section 5 and Annex 4.

# Customer preferences and behaviour

3.28 As explained in Section 2, we have conducted market research to inform our understanding of the customer experience of cloud infrastructure services. Here we provide a summary of some of the key findings from the quantitative phase of our research.

3.29 Since this market study is focused on public cloud infrastructure services, our market research focused on companies and organisations already using IaaS and/or PaaS or actively considering those services. While our research includes some customers who also purchase other services such as SaaS and private cloud, our survey is not attempting to be representative or reflective of these customers.[81]

## Current and expected future cloud use

3.30 In our research the most frequently mentioned reason to adopt cloud computing is greater flexibility and agility.[82] The second most important driver is improved security. These two drivers are also the two most frequently mentioned when respondents were asked to

[80] According to Microsoft this scenario would rarely make economic sense due to significant complexity along a variety of dimensions without any meaningful benefits, such as requiring customers to manage different commitments and capabilities from different providers with respect to security, privacy, regulatory compliance, resiliency, sustainability, procurement and management (and more) for a single cloud workload, Microsoft confidential response to the interim report, page 8. This comment was in response to the interim report definition of integrated multi-cloud – see the interim report, paragraph 3.25(c).

[81] We refer to the results of our research throughout our report. Some of our analysis compares the responses given by subgroups of the total sample of respondents. Due to sample size limitations, our observations are not always based on statistically significant differences. We treat these findings as indicative, and place greater weight and reliance on them where we see consistency with other evidence sources.

[82] Context Consulting research report, slide 29.

identify their single most important reason to adopt cloud computing. Financial reasons appear relatively less important: reduction of capex is mentioned by 22% and lowering overall spend on IT was mentioned by 42% of respondents.

**Figure 3.3: Main drivers to adopt cloud computing, as reported in our market research**

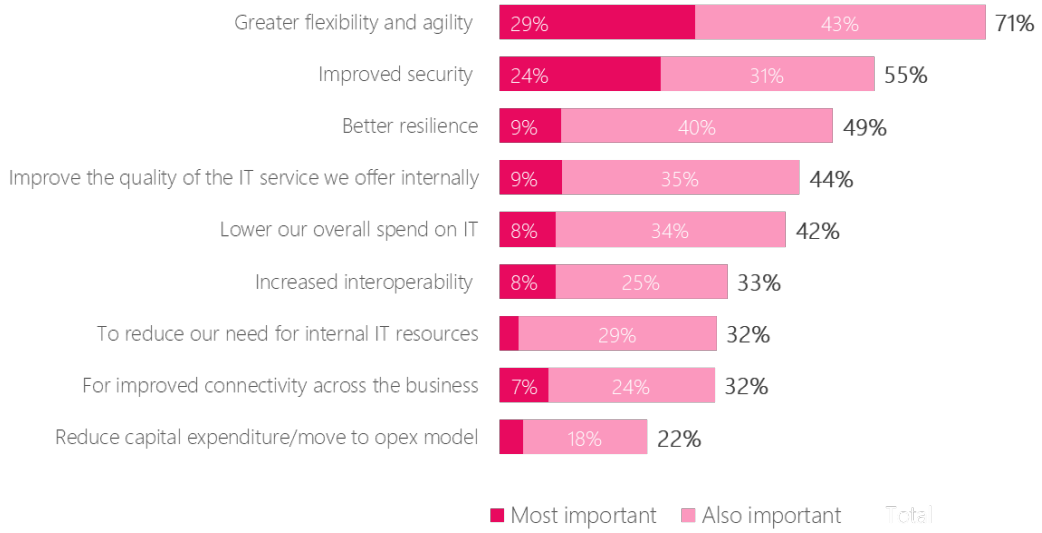


*Source: Context Consulting research report, slide 29.*

3.31 The choice of deployment model (public, private or hybrid), service level (IaaS, PaaS or SaaS), and whether to adopt a multi-cloud model depends on customers' business needs and the level of control they would like to have over their cloud service.

3.32 In our research, we find that a large proportion of respondents purchase more than one type of cloud service, with 33% of respondents purchasing all three (IaaS, PaaS and SaaS) and 35% of respondents purchasing a combination of IaaS and PaaS. Overall, 78% of respondents use SaaS, 69% - IaaS, and 55% - PaaS.[83]

3.33 The three most frequent use cases for current IaaS/PaaS users are storage, databases and back-up (82%, 76% and 71% of respondents, respectively, are currently using cloud for these purposes). IT & Tech companies are more likely than other respondents to be using central processing unit (CPU) - intensive applications (56%) and hosting apps on behalf of clients (53%).

[83] Context Consulting research report, slide 19. We note that given our research focused on companies and organisations already using IaaS and/or PaaS, or actively considering those services, these findings are likely to over-state the use of IaaS and PaaS services relative to users of cloud services as a whole.

**Figure 3.4: Current and potential workloads allocated to cloud, as reported in our market research**

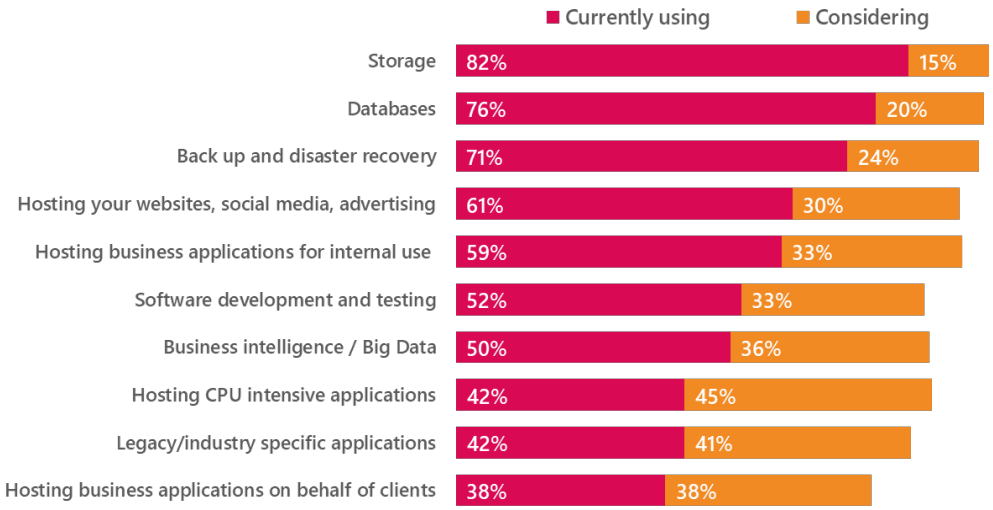


*Source: Context Consulting research report, slide 35.*

3.34 Many respondents have more than one use-case and the average number of use-cases across all respondents is 5.7. 16% of respondents report three use cases or fewer, 31% have 4-5, 31% have 6-7, and 22% have eight to ten. In general, larger organisations, as well as IT & Tech companies, tend to have more cloud use cases.[84] This is consistent with evidence received from the hyperscalers indicating that their customers purchase multiple public cloud services on average, for example (as detailed in Section 5):

a) customers spending more than $10k per year – which account for [✂]% of hyperscalers' revenues – consume at least [✂] [10-20] first-party proprietary products; and
b) customers spending more than $1m a year (accounting for [✂]% of hyperscalers' revenue) on average take at least [✂] [30-40] first-party proprietary services.[85]

3.35 We also asked our research respondents about changes they expected to see in their use of cloud computing in the next 18 months (see Figure 3.5 below). 43% of current users of IaaS/PaaS (or those considering using) reported the intention to migrate more workloads to the cloud in the next 18 months. Attracting IT workers or investing in 'in-house' cloud skills were two further areas where many respondents anticipated making changes. Attracting skilled IT workers is particularly pertinent for healthcare and IT & Tech companies surveyed and was also more important for larger firms compared to smaller ones.[86]

[84] Context Consulting research report, slide 37.
[85] Ofcom analysis of: AWS response dated 14 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' column E); Microsoft response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Confidential Annex 2, tab 'Q1 Customer distribution' column E); Google response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' column E).
[86] "Attracting IT workers …" was mentioned by 45% of respondents with more than 2,500 employees, and by 40% of respondents in the 1,000-2,499 bracket.

**Figure 3.5: Expected changes in cloud computing use, as reported in our market research**

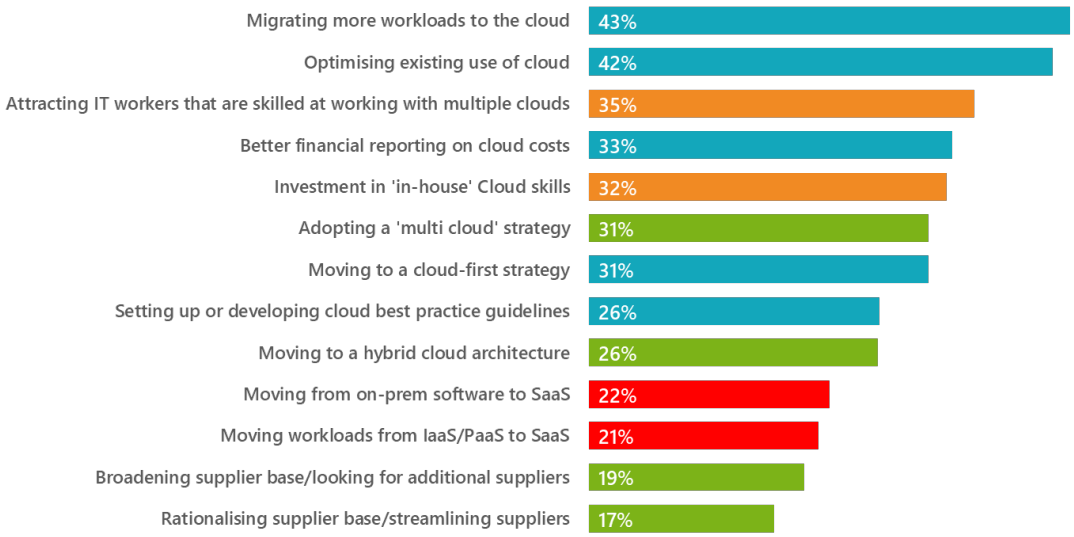


*Source: Context Consulting research report, slide 43.*

*Note: Orange bars represent changes related to staff/skills; green – change in supplier set-up; red – move towards SaaS.*

3.36 For comparison, the Flexera 2023 State of the Cloud report also puts optimisation of existing cloud use, progressing on a cloud first strategy and migrating more workloads to the cloud as top priorities for companies.[87]

3.37 Respondents' focus on migrating more workloads to the cloud also aligns with IDC projections that the share of total IT spending (including IaaS spend and PaaS spend) deployed on cloud services will increase significantly over the next five years, both in the UK and globally.[88]

## Spend on cloud services has been increasing and is expected to continue growing

3.38 For 82% of respondents spend on cloud services has increased in recent years (for 26% it increased greatly, and for 57% slightly). 16% of respondents reported that their spend stayed about the same, and for about 1% it decreased. The increase in spend is consistent across different respondent characteristics, including company size, industry and stage of cloud adoption.

3.39 Among respondents who reported that their spend 'increased greatly' and 'increased slightly', the groups that experienced an increase of cloud spend more often than average are larger companies (88-90% for companies with more than 1,000 employees), early adopters of technology (89%) and those using 3 or more providers (90%).

[87] Flexera, 2023. 2023 State of the Cloud Report, page 29 [accessed 29 September 2023].
[88] IDC, 2023. Worldwide Black Book: Live Edition, July (V2 2023) Forecast (published July 2023). Total IT spending includes the following IDC technology categories from the Black Book publication: Infrastructure, Application Development & Deployment, Applications, System Infrastructure Software, Managed Services, Support Services, Project Oriented Services and Devices.

3.40 When asked how they expected their cloud spend to change in the next 18 months, 79% of respondents said they expected it to increase slightly or greatly, and only 18% expect it to stay about the same. Larger organisations are slightly more likely than smaller organisations to expect to increase spend (see Figure 3.6).

**Figure 3.6: Expectation of change in spend on cloud, as reported in our market research**

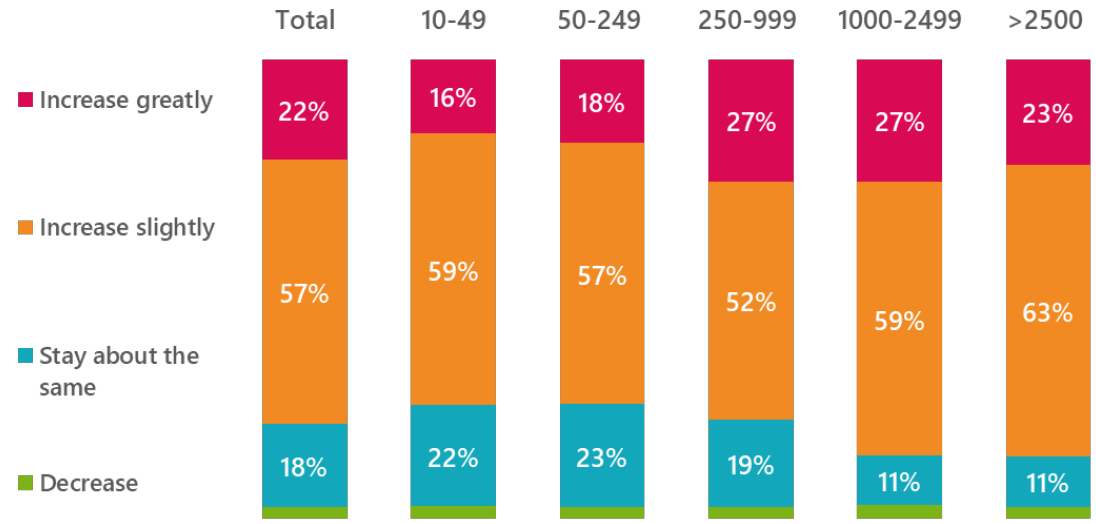


*Source: Context Consulting research report, slide 42.*

## Use of private, public and hybrid cloud

3.41 In our market research, 34% of IaaS/PaaS users said that they only use private cloud, 12% use only public cloud, 44% use both and they are integrated with each other, and 10% use both but the two are part of separate IT architectures.

**Figure 3.7: Use of private, public and hybrid cloud, as reported in our market research**

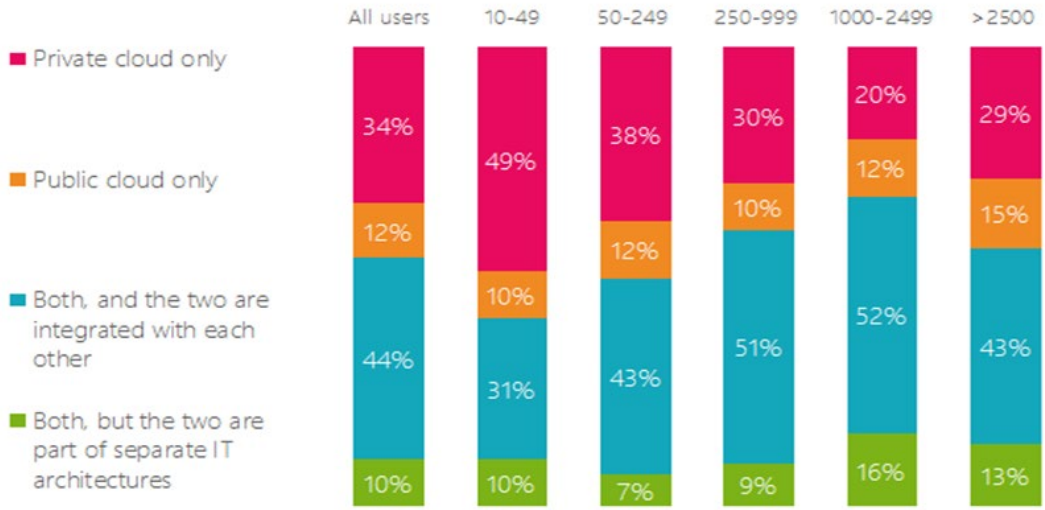


*Source: Context Consulting research data tables, Q21.*

3.42 Smaller companies (with 10-49 employees), public sector organisations, companies that do not use the hyperscalers and companies that use only one provider are more likely to say that they use only private cloud. Younger companies tend to report using 'public cloud only'

more often than average (at 32% compared to 12% overall). Larger companies, those in IT & Tech and companies that use three or more IaaS/PaaS providers tend to use integrated public and private cloud more than others.

3.43 It is possible that some respondents to our market research may have misunderstood the distinctions between private cloud, public cloud and on-premises infrastructure, so the numbers of organisations that reported using 'private cloud only' may be over-stated. In subsequent sections we sometimes provide findings from the research that exclude 'private cloud only' respondents from the base, in addition to providing findings across all respondents.

3.44 In contrast, the Flexera 2023 State of the Cloud report[89] surveyed (mainly large) organisations from across the world who purchased IaaS, PaaS and SaaS – therefore capturing a different customer base to our market research (so the respective results are not directly comparable). Flexera found that 'public only' cloud use (either single or multiple providers) was reported by 24% of their respondents, while 'private only' (either single or multiple providers) was reported by 4% only. Flexera reports that most companies in their sample are taking a hybrid approach, combining the use of both public and private clouds.

## Market players

3.45 The cloud services supply chain is complex, involving different types of suppliers providing services at some or all levels of the cloud stack. This includes:

a) cloud providers, who are usually present at all levels of the cloud stack (IaaS, PaaS and SaaS);[90]
b) ISVs, who usually do not own any physical infrastructure and are present at only one or two levels of the cloud stack (PaaS and/or SaaS); and
c) suppliers of professional services, who provide customers access to cloud services and/or support for using cloud services.

3.46 We consider the role of these suppliers in further detail below. As illustrated in Figure 3.8 below, cloud providers and ISVs can either directly sell to customers, or indirectly sell to customers through suppliers of professional services and/or marketplaces operated by cloud providers. Marketplaces are an online platform, where cloud providers and ISVs can offer services to customers, which run on the underlying infrastructure of the provider offering that marketplace. We consider marketplaces separately in Section 4.

[89] Flexera, 2023. 2023 State of the Cloud Report, page 18 [accessed 29 September 2023]
[90] We note there are some providers of cloud infrastructure that may offer IaaS only, or IaaS and PaaS without SaaS (e.g. OVHcloud). These providers fall within the scope of our definition of cloud provider.

**Figure 3.8: The cloud services value chain**

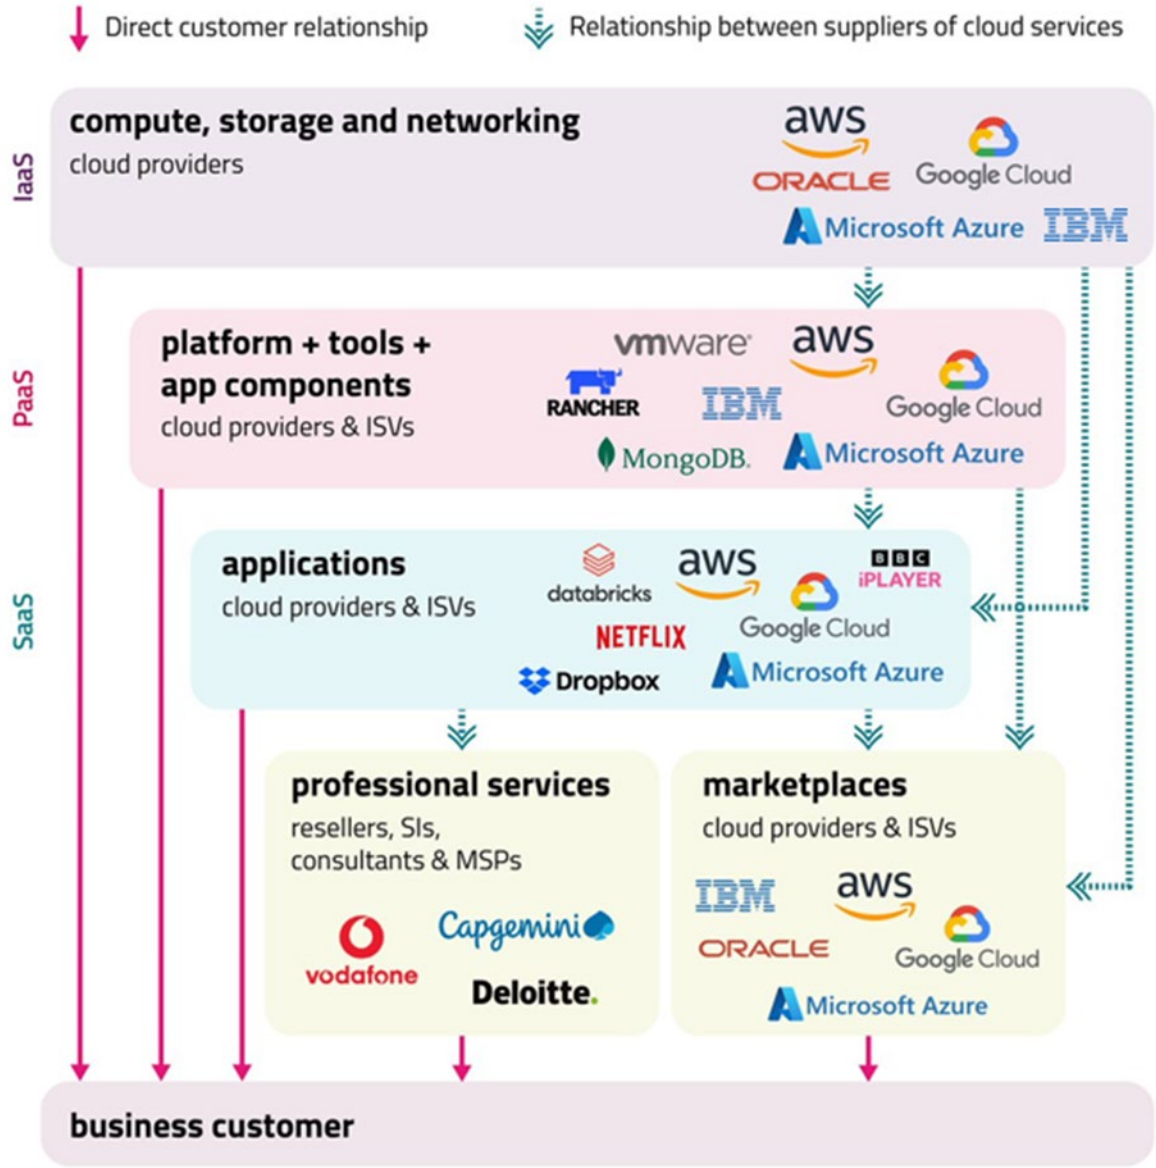


*Source: Ofcom.*

## Cloud providers

3.47 Cloud providers are vertically integrated suppliers of cloud services that operate their own cloud infrastructure (i.e. they own the underlying raw computing resources). They provide the full range of cloud services, in all service (IaaS, PaaS and SaaS) and deployment models. These cloud providers consist of the hyperscalers and a number of smaller providers.

### Hyperscalers

3.48 In the UK there are three hyperscalers - Amazon, Microsoft and Google.[91] They are present at all levels of the cloud stack and provide a wide range of cloud services across multiple product categories at massive scale. Their infrastructure is built on millions of physical servers and virtual machines hosted in huge data centres around the world.

[91] Globally there are other hyperscalers, such as Alibaba, Huawei and Tencent. However, they do not offer cloud services across the cloud stack and at scale in the UK.

3.49 These three hyperscalers are the main suppliers of public cloud infrastructure services in the UK. They collectively account for [✂]% [70% to 80%] of total revenues generated from the supply of public cloud infrastructure services (IaaS and PaaS).[92] They also offer a wide portfolio of SaaS services.

*Amazon*

3.50 Amazon operates through its subsidiary AWS, which is considered as the overall market leader in cloud in the UK. We estimate that, as of 2022, AWS accounted for around [✂]% [30% to 40%] of the UK's public cloud infrastructure revenues.

3.51 Amazon was the first to enter the cloud services market in 2006, after having invested in the relevant IT infrastructure for its own online retail business. AWS started by offering IaaS to customers, with the launch of Amazon Elastic Compute Cloud (EC2), and then subsequently expanded to the PaaS and SaaS layers.[93] AWS currently offers over 210 cloud infrastructure services, that are organised into several product categories (e.g. compute, analytics).[94]

3.52 Amazon, the parent company, also has several other businesses such as Amazon Store (i.e. online retail), Devices and Services, Entertainment (i.e. Prime Video), and Delivery and Logistics.[95] Amazon uses AWS to provide solutions across its businesses, for example, Prime Video uses AWS for compute, database and other services.[96] Amazon categorises its overall operations into three segments: AWS, North America and International.[97] While AWS is Amazon's smallest operating segment by revenue, representing about 16% of Amazon's revenue in the year to December 2022, it was Amazon's only profitable segment in that year and it remains Amazon's most profitable segment as of the quarter to June 2023.[98]

*Microsoft*

3.53 Microsoft is the second largest cloud provider in the UK. We estimate that, as of 2022, Microsoft accounted for around [✂]% [30% to 40%] of the UK's public cloud infrastructure revenues.

3.54 Before entering cloud, Microsoft had already established itself as a major player in the provision of operating systems (through Windows OS) and productivity software (through Microsoft 365). Microsoft first entered the cloud computing market in 2008, with a PaaS offering which enabled developers to deploy applications in the cloud.[99] This offering was later generally made available across countries in 2010. Soon after, Microsoft launched Office 365 (a SaaS level service), before extending its presence to IaaS.[100] Microsoft currently

---

[92] Ofcom analysis of data provided in response to our statutory information requests and data from Synergy and IDC. See Section 4 for more detail on UK shares of supply.
[93] AWS website. About AWS [accessed 28 July 2023].
[94] Ofcom analysis of IaaS and PaaS products listed on AWS's website.
[95] AWS website. What We Do [accessed 28 July 2023].
[96] AWS website. Prime Video Boosts Scale and Resilience [accessed 28 July 2023].
[97] The 'North America' and 'International' segments largely consist of revenues from retail sales of consumer products.
[98] Ofcom analysis based on AWS's published financial statements. Amazon.com announces fourth quarter results [accessed 28 July 2023]. Amazon.com announces second quarter results [accessed 4 August 2023].
[99] Microsoft website. About Microsoft [accessed 28 July 2023]; and Microsoft launches Windows Azure [accessed 28 July 2023].
[100] Microsoft website. Windows Azure Platform Now Generally Available [accessed 6 March 2023 – Paige not found]; Office expands to the cloud [accessed 28 July 2022]; and The History of Microsoft Azure [accessed 28 July 2022].

offers over 200 cloud infrastructure services, across several product categories (e.g. compute, IoT).[101]

3.55 Microsoft organises its services and products into three operating segments: Productivity and Business Processes (e.g. Office 365, LinkedIn), Intelligent Cloud (e.g. cloud services, enterprise services) and More Personal Computing (e.g. Windows operating system, search and news advertising).[102] Microsoft's Intelligent Cloud segment includes Azure which provides cloud infrastructure and platform services (IaaS and PaaS). In the year to June 2023, Intelligent Cloud was Microsoft's largest operating segment, representing 41% of revenues and 43% of profit.[103]

3.56 In the past it has been reported that Microsoft had an objective to use Azure to provide a range of cloud services, such as parts of Office 365 and Bing Search, across its other operating segments.[104] [✂].[105] Microsoft reports aggregate revenues for its full range of cloud services (e.g. Azure, Office 365 Commercial and Dynamics 365) under 'Microsoft Cloud'. In the year to June 2023, Microsoft Cloud accounted for 53% of Microsoft revenue.[106] For the purposes of this market study, our primary focus is Azure, but we also take account of the broader range of cloud services that Microsoft provides where relevant.

*Google*

3.57 Google is the third largest cloud provider in the UK. Relative to AWS and Microsoft, Google is significantly behind in terms of its share of supply for IaaS and PaaS. We estimate that, as of 2022, Google accounted for approximately [✂]% [5% to 10%] of the UK's public cloud infrastructure revenues.

3.58 Before entering the cloud market, Google had already established a leading position in a range of digital markets with its search engine, as well as services such as Gmail and Google Maps. Google first entered the cloud market in 2008, with a preview release of Google App Engine, a platform enabling businesses to develop applications (PaaS).[107] This platform was initially only made available to developers and was later made available as an official fully supported product in 2011.[108] By then, Google had also expanded to IaaS, with the launch of Google Cloud Storage, and subsequently to SaaS. Google currently offers more than 190 cloud infrastructure services, that are organised into several product categories (e.g. compute, data analytics).[109]

3.59 Alphabet, the parent company of Google, organises its operations into three segments: Google Services (e.g. advertising, Google Maps, YouTube), Google Cloud (e.g. Google Cloud Platform, Google Workspace collaboration tools) and Other Bets[110] (combination of all other services).[111] Google Cloud has generated annual operating losses to date (although it

---

[101] Ofcom analysis of IaaS and PaaS products listed on Microsoft Azure's website.
[102] Microsoft website. Segment Information [accessed 28 July 2022].
[103] Ofcom analysis based on Microsoft's published financial statements. Microsoft 2023 10-K [accessed 4 August 2023].
[104] ZDNET, 2021. Microsoft moves closer to running all of its own services on Azure [accessed 28 July 2022].
[105] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[106] Ofcom analysis based on Microsoft's financial statements. Microsoft 2023 10-K [accessed 4 August 2023].
[107] Google website. Google Cloud Platform [accessed 28 July 2023].
[108] Cloud Guru, 2018. The History of Google Cloud Platform [accessed 28 July 2023].
[109] Ofcom analysis of IaaS and PaaS products listed on Google Cloud's website.
[110] According to Alphabet, revenues from Other Bets are generated primarily from the sale of health technology and internet services.
[111] Alphabet website. Alphabet Announces First Quarter 2022 Results [accessed 28 July 2023].

reported operating profits for the first two quarters of 2023),[112] and it represented around 10% of Alphabet's total revenue in 2022, despite it being the fastest-growing segment of Alphabet in that year.[113]

3.60 The Google Cloud segment includes Google Cloud Platform which provides cloud infrastructure and platform services (IaaS and PaaS). Alphabet uses Google Cloud Platform to provide its own services (e.g. Gmail), and increasingly services across its other operating segments too.[114] For the purposes of this market study, our primary focus is the revenue generated from Google Cloud Platform services within Alphabet's Google Cloud segment, but we also take account of the broader range of cloud services that Alphabet provides.

3.61 In response to our interim report, Google reiterated its view that "it is not accurate to describe the market as having only "*three hyperscalers*"", citing that its market shares are comparatively closer to other cloud providers than to AWS and Microsoft. [115] While we recognise that Google is considerably behind Microsoft and AWS in the terms of overall market shares, Google is growing rapidly and leading the chasing pack by some way, and competes closely with Microsoft and AWS in revenues of certain sub-categories, such as [✂].[116] Furthermore, Google is similar to Microsoft and AWS when considering factors such as global data centre count, number of offered products, investment into cloud specific R&D, and numbers of acquisitions. We note that there is a wide industry consensus regarding hyperscaler classification, including recently published Cloud Report by the *Authorite de la concurrence*.[117] Therefore we remain of the view that Google should be considered a hyperscaler for the purpose of our analysis.

## Smaller cloud providers

3.62 In addition to the hyperscalers, there are also a range of mid-scale (e.g. IBM, Oracle) and small-scale (e.g. OVHcloud, Scaleway) cloud providers offering IaaS and PaaS products that we collectively refer to as 'smaller cloud providers'. In response to our interim report AWS and Microsoft suggested that competition to provide cloud services is increasing and listed a number of companies who have begun offering cloud services.[118] We acknowledge there has been some entry in the last decade, but the scale of these cloud providers remains relatively small.

3.63 In 2022, the smaller cloud providers represented around [✂]% [0-5%] of revenues associated with the supply of public cloud infrastructure services in the UK.[119] These providers compete with the hyperscalers by offering services across the cloud stack, as well

---

[112] Alphabet Inc., 2023. 2023 Q1 10-Q, page 29 [accessed 17 July 2023]; Alphabet Inc., 2023. 2023 Q2 10-Q, page 32 [accessed 27 July 2023].
[113] Ofcom analysis based on Alphabet's published financial statements. 2022 Annual report [accessed 28 July 2023].
[114] Google's website. What is Cloud Computing? [accessed 28 July 2022]. DCD, 2021. Google to migrate parts of YouTube to Google Cloud [accessed 28 July 2023].
[115] Google response to interim report, paragraph 4.
[116] Ofcom analysis of [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; and, [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[117] *Authorite de la concurrence* website. The Autorité de la concurrence issues its market study on competition in the cloud sector [accessed 29 September 2023]
[118] AWS response to the interim report, paragraph 5. Microsoft response to the interim report, paragraph 45.
[119] Ofcom analysis of data provided in response to our statutory information requests and data from Synergy and IDC. See Section 4 for more detail on UK shares of supply.

as partner with the hyperscalers by offering complementary services. Though a number of smaller cloud providers are present in the UK, based on IDC data, no new significant cloud providers offering a broad range of services have entered the UK market since 2018.[120] We illustrate further the role of smaller cloud providers, using IBM and Oracle as examples, below.

### IBM

IBM first started providing cloud services around 2008, with the launch of a SaaS offering. It later expanded to provide IaaS and PaaS.[121] Before entering the cloud market, IBM had already established itself as one of the world's largest IT companies, in producing and selling computer hardware and software. IBM have an extensive history in compute infrastructure.[122]

IBM currently offers a range of services (e.g. compute, AI and machine learning).[123] IBM's strategy has a particular focus on hybrid cloud and multi-cloud solutions. Its capabilities in this area strengthened with the acquisition of RedHat in 2019, an open-source software provider that delivers hybrid cloud technologies.[124] IBM's customers tend to be mainly large and mid-size enterprises.[125] IBM also appears to have a focus on delivering cloud services to regulated industries, for example the financial services sector, by meeting their specific cloud needs for regulatory compliance.[126]

### Oracle

Oracle entered the cloud services market with Oracle Cloud Infrastructure (OCI) in 2016, providing compute, storage and networking services to begin with. It later expanded to provide a wide range of services across IaaS, PaaS and SaaS. Prior to entering the cloud market, Oracle was well known for its on-premises database management systems.[127]

Oracle appear to have a focus on enterprise customers, hybrid and multi-cloud offerings.[128] Oracle has agreements with Microsoft, which allow its services, including its database offerings, to run on Microsoft's platform with "seamless interoperability".[129] Oracle is also recognised to be innovating to meet sovereign cloud needs (the need for data to adhere to regulations of the country where the customer is located).[130] Oracle identify its cloud to be

[120] IDC, Public Cloud Services Tracker 2022 H2 (published April 2023). This data includes ten categories of services across both IaaS and PaaS. Since 2018, the first year in our data, there have been no new companies offering services in five or more of these ten categories. While IDC does not track every company, its dataset includes a broad range of companies, so we would expect any company not separately tracked to be very small.
[121] Datamation, 2017. What is IBM Cloud? [accessed 28 July 2023].
[122] Tech Monitor, 2022. What is IBM known for? [accessed 28 July2023].
[123] IBM website. IBM Cloud products [accessed 28 July 2023].
[124] RedHat, 2019. Press release. IBM Closes Landmark Acquisition of Red Hat [accessed 28 July 2023].
[125] [✂].
[126] IBM website. Regulated workloads [accessed 28 July 2023].
[127] Oracle website. About Oracle and Oracle Cloud Infrastructure Platform [accessed 28 July 2023].
[128] Oracle website. Oracle's distinct approach on hybrid and multi-cloud [accessed 28 July 2023].
[129] Oracle website. Oracle Interconnect for Azure [accessed 28 July 2023].
[130] Techzine, 11 July 2022. Oracle expands OCI with sovereign cloud regions inside EU [accessed 28 July 2023].

"the first and only sovereign, dedicated dual-region cloud for UK Government and Defence customers".[131]

## Independent software vendors (ISVs)

3.64 ISVs are suppliers of cloud services, typically PaaS and/or SaaS, that do not usually own the underlying infrastructure.[132] ISVs and cloud providers interact in a number of different ways:

a) ISVs rely on cloud providers as input suppliers, i.e. they may use IaaS from cloud providers to develop one or more downstream services, such as PaaS offerings (e.g. Databricks, MongoDB, OpenShift, Twilio, VMware) and/or SaaS offerings (e.g. Atlassian, Blue Prism, Snowflake);[133]
b) ISVs can also compete directly with cloud providers, offering services at the same layer of the cloud stack. In other circumstances they may complement the services of cloud providers, including the hyperscalers; and
c) ISVs may further rely on cloud providers, and the hyperscalers in particular, as distributors. This could be through cloud providers directly selling ISVs' services, offering ISVs a platform through which to sell their services (such as a marketplace), or by offering ISVs access to customers. We refer to this in Section 4.

3.65 There are many ISVs, with different specialisms, present in the UK cloud services market. ISVs that provide PaaS tend to compete in specific product categories, rather than across PaaS as a whole. We are taking account of evidence regarding, and from, many ISVs within this market study. To provide some insight into the role of ISVs, we summarise the roles of a small sample of different types of ISVs as illustrative examples below.

### VMware

VMware was founded in 1998. It specialises in providing virtualisation technology, having first developed the technology for on-premises use. This technology allowed users to run multiple operating systems, as virtual machines, on a single physical machine. In 2009, VMware launched its virtualisation technology for cloud computing.[134] VMware's strategy is focused on enabling multi-cloud environments.[135]

VMware partners with a range of cloud providers. For example, it delivers its own cloud services on AWS, by providing customers its virtualisation products on AWS's cloud as well as access to AWS's native services (e.g. EC2).[136] It also partners with a range of other cloud providers (e.g. Microsoft, Google, IBM) enabling them to create cloud offerings using VMware's solutions on their own infrastructure. At the same time, VMware also appears to compete with cloud providers' own cloud native virtualisation solutions.[137]

---

[131] Oracle website. Oracle Cloud for UK Government & Defence [accessed 28 July 2023].
[132] Although, we note that there may be exceptions to this. For example, Salesforce is an ISV that also operates its own infrastructure.
[133] We recognise that the PaaS/SaaS categorisation can be arbitrary.
[134] VMware website. VMware Timeline [accessed 28 July 2023].
[135] VMware website. Multi-Cloud Environments [accessed 28 July 2023].
[136] AWS website. VMware Cloud on AWS [accessed 28 July 2023].
[137] VMware response dated 9 December 2022 to the s.174 notice dated 27 October 2022, Part A, question 14, paragraph 14.2.

## MongoDB

MongoDB was founded in 2007. It specialises in database management and document databases.[138] MongoDB Atlas, its PaaS cloud offering, is positioned as a developer data platform. It gives developers the ability to run their databases across several cloud providers and provides them with access to a range of features and tools, enabling users to access, query and analyse data.[139]

Atlas can currently be deployed on the infrastructure of AWS, Microsoft and Google. Atlas is built in a way that makes it complementary to the customers cloud provider's infrastructure. [✂].[140] AWS and Microsoft have their own document databases, referred to as DocumentDB and CosmosDB respectively, however the capabilities of these databases are thought to be different to MongoDB Atlas.[141]

## Snowflake

Snowflake was founded in 2012.[142] It specialises in providing data warehouses (SaaS offering) which provides users with the ability to store and access structured and unstructured data. Snowflake also offer a 'cloud data platform' which is capable of supporting multiple data workloads from data warehousing to data engineering, across several cloud providers.[143]

Snowflake's offerings run on top of the infrastructure provided by public clouds. Snowflake's offerings were initially only available on AWS, and later also became available on Azure and Google Cloud.[144] Snowflake also identify the hyperscalers as its competitors.[145] Each of the hyperscalers have a cloud offering that can be considered comparable to some extent to Snowflake's offerings (i.e. Amazon RedShift, Google BigQuery, Azure SQL Data Warehouse).

## Yugabyte

Yugabyte was founded in 2016. It specialises in providing database technology. The YugabyteDB offering provides customers with open-source distributed databases (typically used to store data across multiple sites), and access to enterprise database features.[146] YugabyteDB is available to customers as both fully managed and self-managed services.[147]

YugabyteDB can run on top of the infrastructure provided by public and private clouds. Yugabyte identify themselves to be customers, partners and competitors of all three hyperscalers. In that, they purchase infrastructure services from the hyperscalers, partner

---

[138] MongoDB website. Our Mission [accessed 28 July 2023].
[139] MongoDB website. MongoDB Atlas [accessed 28 July 2023]; and Advantages of MongoDB [accessed 28 July 2023].
[140] [✂].
[141] [✂].
[142] Techstory, 2020. The Snowflake story [accessed 28 July 2023].
[143] Snowflake website. Enabling the data cloud with Snowflake [accessed 28 July 2023].
[144] Snowflake website. Cloud partners [accessed 28 July 2023].
[145] Snowflake response dated 7 December 2022 to the s.174 notice dated 26 October 2022, question 12, page 11.
[146] Yugabyte website. About Yugabyte [accessed 28 July 2023].
[147] Yugabyte website. YugabyteDB Deployment Options [accessed 28 July 2023].

with hyperscalers to win clients, and also compete with the database offerings of hyperscalers.[148]

## Suppliers of professional services

3.66 Suppliers of professional services provide customers access to cloud services and/or support to using cloud services. These suppliers can take several different roles. For example, they could act as system integrators, managed service providers, consultants and/or resellers.[149] Many of the largest suppliers of professional services (such as Accenture, Capgemini, Cognizant and Deloitte) take on multiple of these roles. We set out some examples below.

a) **System Integrators (SIs)** bring together the different cloud solutions a customer purchases. Examples include: Accenture, Capgemini, Deloitte, Vodafone.
b) **Managed service providers (MSPs)** set up customers, so that they can run their operations on the cloud infrastructure of cloud providers, whilst offering long-term managed services (service management, compliance support, etc.). Examples include: Cloudreach, Infosys, Logicworks, Wipro.
c) **Consultants** provide advice to customers on their use of cloud, such as their choice of supplier, initial cloud migration, multi-cloud strategies, etc. Examples include: Accenture, Capgemini, Cognizant, Deloitte.
d) **Resellers** essentially resell cloud services from cloud providers. They may also provide some value-added services on top and/or tailor the pricing models of cloud providers. Examples include: Insight, Strategic Blue.

3.67 Customers typically choose to use these suppliers, rather than purchase directly from cloud providers, where there is a need for further expertise or support with their cloud purchases. These suppliers can add value in a range of ways. For example, they may offer customers additional or bespoke services that build on cloud providers' products (e.g. SIs/MSPs); provide access to lower prices (e.g. resellers); and/or provide technical expertise to customers on aspects such as cloud migration (e.g. consultants).

3.68 We observe that most suppliers (e.g. Accenture,[150] Capgemini[151]) tend to work with cloud providers and ISVs by forming partnerships. For example, we understand from Capgemini that it is common for them to work with the hyperscalers (and other smaller cloud providers) to win procurements, as well as to build industry-specific solutions which can then be proposed to clients.[152]

3.69 We are also aware of cases where some suppliers indicate that they compete with the hyperscalers. For example, [✂].[153] However, we note such cases appear overall to be limited.

[148] Ofcom / Yugabyte meeting, 22 November 2022.
[149] This is Ofcom's view of the different types of suppliers of professional services. We note that these categorisations are not necessarily distinct and other categorisations may be possible.
[150] Accenture website. Cloud services [accessed 28 July 2023].
[151] Capgemini website. Technology partners [accessed 28 July2023].
[152] Ofcom / Capgemini meeting, 1 November 2022.
[153] Ofcom / [✂] meeting, [✂].

# Cloud services in the telecoms and broadcasting sectors

## Telecoms

3.70 Cloud services are changing how telecoms services are being produced and delivered to customers and so it is important to understand how the telecoms sector is using cloud services.

3.71 While there are differences between telecoms providers and their usage of cloud, we can broadly group telecoms cloud usage into three categories: telecoms network functions, multi-access edge computing (MEC) and enterprise IT functions.

### Telecoms network functions

3.72 UK telecoms providers told us that some of their network functions[154] have started moving onto the cloud. Where telecoms providers are using the cloud to host these network functions, they are broadly using private cloud.[155] BT Group told us it thinks telecoms providers who have invested in their private clouds are unlikely to migrate fully to the public cloud,[156] suggesting hybrid cloud architecture will be used in the future. This view was supported by [✂], [✂] and [✂].[157] [✂] also told us that they have plans to begin testing of some 'non-critical' network functions in the public cloud.[158]

3.73 Telecoms providers highlighted some drawbacks to using public cloud for telecoms specific workloads including a lack of control, a risk to resiliency and uncertain costs.[159] For example, in their CFI response, BT Group noted the primary reason they do not use public cloud services within their UK core network is security and resilience, as well as the desire to retain end-to-end control of their network assets.[160]

3.74 Despite this, our conversations with telecoms providers made clear that they recognise a number of key benefits that the public cloud offers. A telecoms provider [✂] noted the speed of innovation, automation and implementation of updates that public cloud providers offer; the flexibility in scaling up and down, which could enable telecoms providers to be more reactive and agile to customer needs; and the opportunity to save money by reducing their expenditure on in-house data centres and resourcing.[161]

[154] Network function is a component of telecom networks that delivers a specific function (e.g. router, switch, load balancer, firewall).
[155] Capgemini, 2023. Networks on Cloud: a clear advantage, page 6 [accessed 28 July 2023].
[156] BT Group response to the CFI, page 6, paragraph 2.
[157] Ofcom / [✂] meeting, [✂], Ofcom / [✂] meeting, [✂] and Ofcom / [✂] meeting, [✂].
[158] Ofcom / [✂] meeting, [✂] and [✂] response dated [✂] to our proposed use of information dated [✂].
[159] Ofcom / [✂] meeting, [✂] and Ofcom / [✂] meeting, [✂].
[160] BT Group response to the CFI, page 6, paragraph 1.
[161] Ofcom / [✂] meeting, [✂].

3.75 We have seen recent examples of new cloud products being developed and launched specifically for telecoms providers.[162] As the suite of products targeted at telecoms providers grows, we may see increased use of public cloud services for some network functions.[163]

3.76 Research from global technology intelligence firm ABI research, suggests that while the private cloud is currently the preferred approach for most telecoms providers, some mobile operators outside the UK are already using or planning to use public cloud for network functions.[164]

### Multi-access edge computing (MEC)

3.77 MEC allows the processing of workloads and storing of data close to the edge of a telecoms network, i.e. the physical location where users connect with the telecoms network. Some telecoms providers have already begun to offer MEC, including Vodafone through a partnership with AWS.[165] BT Group have also recently announced a partnership with AWS to deliver MEC services, with the first site already live for customer trials, and general availability targeted for later this year.[166] Other providers told us that they are interested in MEC and continue internal considerations regarding expansion into this area, so we may see an increase in MEC offerings from telecoms providers in the future.[167]

3.78 A survey of mobile operators conducted by researchers Heavy Reading suggests that the main driver for operators to move to edge computing was more efficient use of bandwidth.[168] Larger operators tend to be driven by the opportunity to open new revenue streams through differentiated services (such as video conferencing) against competitors, whereas smaller operators tend to be focused on improving application performance and resilience.

### Enterprise IT functions

3.79 Like many other large businesses, telecoms providers are using cloud services to some extent in the delivery of many of their internal business and IT functions.[169] UK telecoms providers use public cloud services largely provided by the three hyperscalers for their enterprise IT functions. They use, or have the intention to use, multiple providers as part of a multi-cloud strategy with distinct workloads (siloed multi-cloud). The drivers for moving to

---

[162] Examples include AWS for Telecom | Industry Cloud Solutions | AWS (amazon.com); Telco Network Automation – AWS Telco Network Builder – Amazon Web Services; Telecommunications Industry Solutions | Microsoft Industry; Azure for Operators – Telecom Solutions | Microsoft Azure; Telecommunications industry tools|Google Cloud; IBM Telecommunications solutions; Cloud Infrastructure for Communications | Oracle United Kingdom; Telco Cloud Platform | VMware and Telecom Data Cloud | Snowflake for Telecommunications [accessed 24 August 2023].

[163] There are some examples of operators around the world who have plans to use public cloud for network functions. Examples include, Rakuten Mobile (Japan), Dish (USA), AT&T (USA) and O2 Telefónica (Germany). Capgemini, 2023. Networks on Cloud: a clear advantage [accessed 28 July 2023].

[164] ABI research, 2023 Whitepaper | Public, Private, or Hybrid Cloud for 5G Core Network Deployments (abiresearch.com) [accessed 28 July 2023].

[165] Vodafone website. Multi-Access Edge Computing [accessed 28 July 2023].

[166] BT Wholesale website. BT uses Amazon Web Services (AWS) Wavelength. [accessed 28 July 2023].

[167] Ofcom / [✂] meeting, [✂]; and Ofcom / [✂] meeting, [✂].

[168] Heavy Reading, 2022. 5G Network Strategies: Operator Survey 2022 [accessed 28 July 2023].

[169] BT Group response to CFI, page 7, paragraph 2, [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

public cloud and choosing a single or multiple cloud providers are similar to those of other large business customers.

### Telecoms providers and the cloud services supply chain

3.80 Telecoms providers sometimes have a multi-layered relationship with the hyperscalers. In addition to being customers they are, in some instances, partners of hyperscalers.[170] Telecoms providers also sometimes act as intermediaries for the sale of cloud services – our market research found that some IaaS and PaaS users purchase cloud services through telecoms providers (see Table 4.1 in Section 4).

3.81 In some other circumstances telecoms providers regard cloud providers as competitors, where they are using cloud technology to offer localised connectivity directly to customers and to transfer large files across the UK and the world.[171]

## Broadcasting

3.82 UK broadcasters are heavy users of public cloud services across the cloud stack, which they use across both their internal operational systems and broadcasting-specific functions, including:

a) content production, for example to set up remote artist workstations, editorial workflows, collaboration platforms and rendering; and
b) content distribution, for example to underpin BVoD services, e.g. BBC iPlayer and ITVX, and content distribution networks (CDNs).

3.83 We received evidence that some broadcasters use a combination of public and private cloud offerings. [✂] told us that key drivers for using private cloud include low latency and information security policy requirements.[172] However, [✂] noted that they do not use private cloud, as they find public cloud offerings to be sufficient in areas such as security and resilience.[173]

3.84 Broadcasters have reported that the main driver for choosing a particular cloud provider is an understanding of broadcasting needs, and the availability of broadcasting-specific tools and apps on the provider's platform. They also require the ability to operate at scale and support the vast amounts of data storage and processing. Production and storage of audio-visual content requires a lot of data and broadcasters have therefore favoured the three hyperscalers, who operate at the requisite scale. Precisely because of these data requirements, some broadcasters have kept significant on-premises IT to manage costs.

3.85 [✂] told us that broadcasters need the most up-to-date, audience friendly digital products available to ensure they can confidently compete in the market. They said that this places limitations on where broadcasters can get their cloud services, as they need providers to support the full range of services that enable streaming.[174]

---

[170] For example, Vodafone's partnership with AWS: Vodafone, 2021. Press release. Vodafone uses AWS Wavelength to launch first Multi-access Edge Computing services in European region and BT Group's partnership with AWS: BT Group, 2023. Press release. BT Group expands partnership with AWS: new deal targets $500m opportunity in connectivity and digital services [accessed 24 August 2023].
[171] AWS website. AWS Private 5G [accessed 28 July 2023].
[172] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[173] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[174] [✂] response dated [✂] to the CFI, page [✂], paragraph [✂].

3.86 From responses to our CFI and statutory information requests, it appears that AWS has been particularly successful in winning and retaining customers in the broadcasting sector. A key reason for this has been the technical expertise of AWS in streaming and content production, which it gained in particular through its acquisition of Elemental in 2015. However, as discussed in Section 5 and 6, there may be certain barriers to switching and multi-cloud that have contributed to AWS's position with broadcasters.

## Cloud services in the public sector

3.87 Public sector organisations represent a significant source of revenue for cloud providers, so competition for these customers is an important aspect of the market. As with many other sectors, cloud services are an increasingly important input into public sector organisations.[175] The UK Government operates a Cloud First policy, where public sector organisations must consider public cloud solutions first before other deployment options.[176]

3.88 To support government departments with this procurement, there is a range of guidance, frameworks, and agreements in place with a range of suppliers to ensure the best possible outcomes.[177] Our understanding is that the processes and guidance designed by the Crown Commercial Service have the objectives of:

a) Securing appropriate services that meet the needs of the departments.
b) Mitigating the risk of lock-in as much as possible.
c) Promoting competition in the public sector marketplace.
d) Securing value for money with any procurement activity.

3.89 One of the procurement routes in which government departments access cloud services is via a digital marketplace, which hosts has 31,000 products across different providers and layers of the stack. In 2022-23, the G-Cloud framework was used to purchase £363m of hosting services.[178] Most contracts typically last two years – a design choice to enable regular competition for these contracts. The government has also signed Memorandums of Understanding which enable a cross-governmental discount for purchases with certain providers, plus additional benefits such as free cloud skills training for agreed numbers of civil servants.[179]

[175] When we mean public sector in this context, we are principally referring to central government departments and arm's length bodies.
[176] Central Digital and Data Office, 2023. Government Cloud First policy [accessed 12 September 2023]
[177] Central Digital and Data Office, 2023. Government Cloud First policy [accessed 12 September 2023]; Central Digital and Data Office, 2020. Creating and implementing a cloud hosting strategy [accessed 12 September 2023].
[178] G-Cloud is an online catalogue where public sector customers can buy cloud-based computing services such as hosting, software, and cloud support. It includes many off-the-shelf, PAYG cloud solutions (Crown Commercial Service, G-Cloud [accessed 18 September2023]); Crown Commercial Service, G-Cloud and DOS Sales [accessed 04 August 2023].
[179] Crown Commercial Service and Central Digital and Data Office, 2023. Cross-departmental Memorandums of Understanding [accessed 12 September 2023]; Microsoft, 2021. UK government signs new three year memorandum of understanding with Microsoft [accessed 12 September 2023]; Amazon, 2020. One government value agreement accelerating cloud adoption innovation across UK government [accessed 12 September 2023].

3.90 In response to our consultation, some respondents noted the risk of hyperscalers dominating the public sector part of the cloud infrastructure services market.[180] Two individual respondents went further, citing concerns about a lack of competition in the existing processes, and the impacts this could have across resilience, security of personal data, the UK economy.[181]

3.91 We recognise these are relevant risks related to a lack of competition, and we note that whilst we have not conducted a complete market assessment for the public sector, the leading positions of AWS and Microsoft on G-Cloud Framework appear similar to our findings for the wider market.[182]

[180] Name withheld 1 response to the interim report, paragraph 4.2. We note the figure cited in this answer is only in relation to G-Cloud spend, which is not the complete picture of government spend on public cloud services. Oracle, response to the MIR consultation, page 3. While this is a US example, we consider the concerns to be relevant to the UK.

[181] Name withheld 1 response to the interim report, paragraphs 5.2 and 6.1; Name withheld 2 response to the MIR consultation, paragraphs 2.2 and 3.2.

[182] Crown Commercial Service, G-Cloud and DOS Sales [accessed 04 August 2023].

# 4.How competition works

## Introduction

4.1 This section looks at how competition works in the cloud infrastructure services market. First, we outline how customers buy cloud services and then assess the factors customers consider when choosing between providers of public cloud services. Second, we consider how suppliers compete for customers across a number of dimensions, including price and quality. We consider how cloud services are sold and the role of ecosystems in the market. Third, we consider key market outcomes, including the positioning of the leading providers. We look at the market shares of UK cloud services, analyse the profitability of the key providers, and consider the service offerings and capabilities of each hyperscaler.

4.2 Respondents to our interim report generally agreed with our assessment of how competition works and market outcomes.[183] Where there were specific comments or disagreements with our assessment, these are highlighted throughout the section.

## How customers choose and buy cloud services

4.3 In this subsection we outline how customers buy cloud services including the channels they use, the purchase process, and the factors customers consider when choosing between providers of cloud services. We draw on our market research, external reports by market analysts, and conversations with cloud customers.

4.4 Customer preferences may vary and the weight they place on each factor will differ according to their needs and access to technical skills. Therefore, we start our assessment by setting out a number of dimensions to categorise customers against, which we then refer to in our later analysis.

### Customer types

4.5 There is no uniform categorisation of customers that cloud providers use in their normal course of business. Some providers loosely classify customers by size (often based on number of employees), and others by industry. Some separate private and public sector clients, whereas others do not. Industry reports, such as Flexera 2023 State of the Cloud Report,[184] Oracle 2023 Report,[185] and others also often classify cloud customers by industry and size, as well as geography. Responses to the interim report did not provide any additional categorisation of customers.

---

[183] University of East Anglia response to the interim report, page 3; Federation of Communication Services response to the interim report, page 3; Sustainable Digital Infrastructure Alliance response to the interim report, page 1; [✂] response to the interim report, page [✂]; [✂] response to the interim report, page [✂]; [✂] response to the interim report, page [✂]; [✂] response to the interim report, page [✂]; [✂] response to the interim report, page [✂]; Google response to the interim report and to the MIR consultation, page 1.

[184] Flexera, 2023. 2023 State of the Cloud Report [accessed 26 September 2023].

[185] S&P Global Market Intelligence (commissioned by Oracle), 2023. Multicloud in the Mainstream: Making IT Work 'As Advertised'. Commissioned by Oracle. S&P Global Market Intelligence [accessed 26 September 2023].

4.6 Depending on the context, different characteristics can be relevant when customers decide what cloud products to purchase and which providers to purchase from:

a) **Size** – larger companies may have more use cases for cloud services, buy larger volumes of cloud services, and have more in house IT capability. Our analysis suggests that cloud providers tend to have a number of large, high-spend customers (e.g. enterprises or government departments) that account for a very small proportion of their customer base but a large proportion of their revenues.[186] Approximately the top 1% of customers account for the majority of revenues for each of the leading cloud providers we looked at. [187] A greater number of small and medium sized businesses represent the majority of their customer base but account for a small proportion of their total revenues.[188]
b) **Industry** – some industries have specialised use cases due to regulatory requirements (e.g. financial sector), or prescribed procurement rules (public sector).
c) **Complexity** – some companies will use cloud to build their own bespoke applications, which can come with increased complexity of needs particularly where solutions require close integration between several distinct cloud services.
d) **Stage of cloud adoption** – companies in the early stage of cloud adoption are more likely to concentrate on migrating workloads into cloud, whereas those in later stages will be working on optimising cloud use. In our market research, for example, companies that have been using cloud longer are more likely to have a formal cloud strategy, and are more likely to use a multi-cloud strategy to mitigate the risk of lock in. Those in later stages may also seek to change suppliers – switch some or all components of their applications to a different cloud (i.e. switching between clouds); or, to an ISV hosted on the same cloud (i.e. switching within clouds).
e) **Approach to technology adoption overall** – in our market research there are often differences in responses between those our research agency has classified as "early technology adopters" and "laggards".[189] For example, early adopters are more likely to have a formal cloud strategy, are more likely to use more than one provider, and are more likely to have switched in the past or to have taken on an additional provider.

4.7 Where relevant, we will highlight differences in our analysis between the various customer characteristics.

# How customers buy cloud services

## Purchase channels

4.8 Our market research found that there are various purchase channels for cloud services. About a half of IaaS users purchased those cloud services from a hyperscaler directly, and

[186] Ofcom analysis of data from: [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[187] Ofcom analysis of data from: [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]. This figure is disputed by one hyperscaler. [✂]. ([✂] response dated [✂] to our proposed use of information dated [✂]).
[188] [✂] response dated [✂] to [✂] of our follow-up email dated [✂] concerning the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to [✂] of our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[189] Context Consulting research report, slide 20.

about 40% of PaaS users did so (in combination with a third party, or just purchasing via a hyperscaler directly). Among IaaS users, we found this is more likely among larger companies (60%), IT & Technology (57%), early adopters of technology (63%), those using 3 or more providers (60%), and those who have switched providers in the past (62%).

**Table 4.1: Channels of purchase of IaaS and PaaS, as reported in our market research**

| Channels of purchase[190] | IaaS users | PaaS users |
|---|---|---|
| **Direct from hyperscaler** | 51% | 38% |
| **Via a telecoms provider** | 33% | 36% |
| **Via a service integrator** | 32% | 30% |
| **Via a managed service provider** | 32% | 36% |
| **Via other provider** | 0% | 1% |

*Source: Context Consulting research data tables, Q33.*

4.9 For PaaS, the following groups are **more** likely to go **via a direct channel** (38% overall): larger companies (44% for those with 1,000-2,499 employees and 47% for >2,500 employees), early adopters of technology (45%) and those who had not considered switching (44%).

4.10 We asked the respondents why they were using a third party or were buying directly, and the main reason for their approach. The most frequently cited reasons for purchasing **directly** were to get a better price/deal, and to get better advice/expertise, at 43% and 39% respectively. Customer support during the purchase and after the purchase are mentioned by 35% of respondents each.

**Figure 4.2: Reasons for purchasing IaaS and PaaS directly, as reported in our market research**

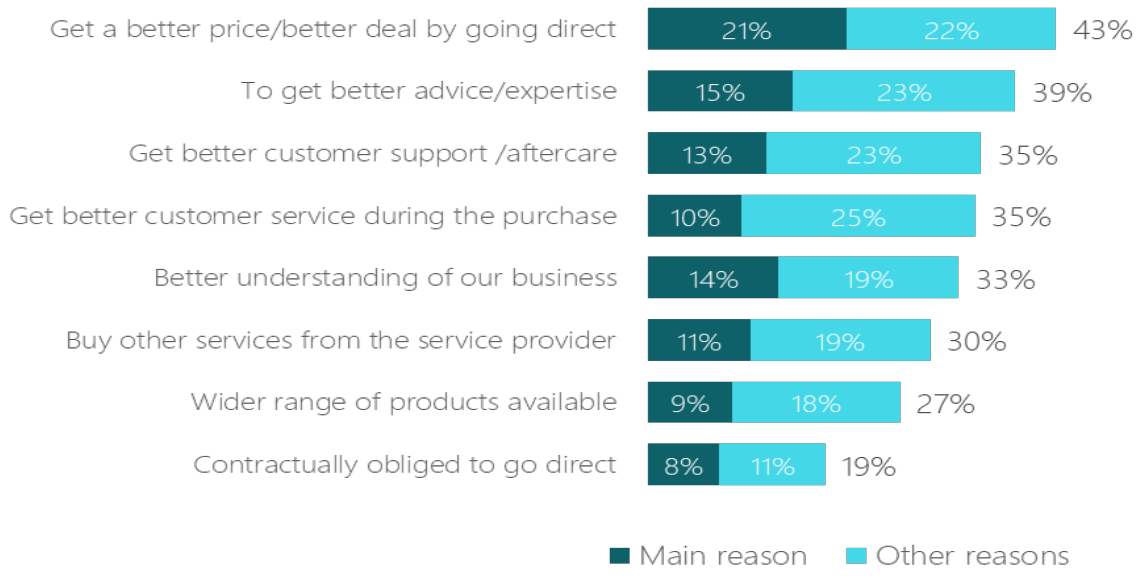


*Source: Context Consulting research report, slide 87.*

[190] These responses are not mutually exclusive.

4.11 The main reasons for going via a third party also include better advice (37%) and customer support (35%), alongside a better understanding of the respondents' business (30%) while the price factor is further down on the list of main reasons (at 29%).[191] As the reasons given for going direct or via a third party are similar, this might suggest that customer characteristics are a better indicator of what is driving this choice. For example, for larger customers, going direct may be the best way to get better advice/expertise but for smaller customers going via a third party may be the best way to get advice.

4.12 Many of the large customers that responded to our statutory information requests and customer questionnaire told us they are not using intermediaries ([✂]).[192] Those that have used intermediaries in the past, did so to get consulting services, or in the areas where they felt the intermediary could provide a greater level of expertise, to facilitate billing, or as a vehicle for the transaction.

4.13 [✂] told us that in their experience, the hyperscalers are geared up to help with 'greenfield' deployments (i.e. where there is no existing cloud infrastructure) that follow a well-established template. In contrast, larger existing enterprises such as [✂], that need to integrate cloud providers with an existing IT footprint, are less well supported directly by the hyperscalers, and these projects can be supported better by intermediaries.[193]

*Marketplace use*

4.14 Of those going directly via a cloud provider, one way to do so is via the marketplaces they manage offering cloud services. Overall, 51% of users in our market research reported using marketplaces. Compared to an average respondent, the following categories were using marketplaces **more** often: early adopters of technology (74%), those using 3+ providers (71%), those who switched in the past (82%), and those who added a provider (73%).[194]

[191] Context Consulting research report, slide 87.
[192] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]. [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂].
[193] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[194] Context Consulting research data tables, Q45.

**Figure 4.3: Reasons to use marketplaces, as reported in our market research**

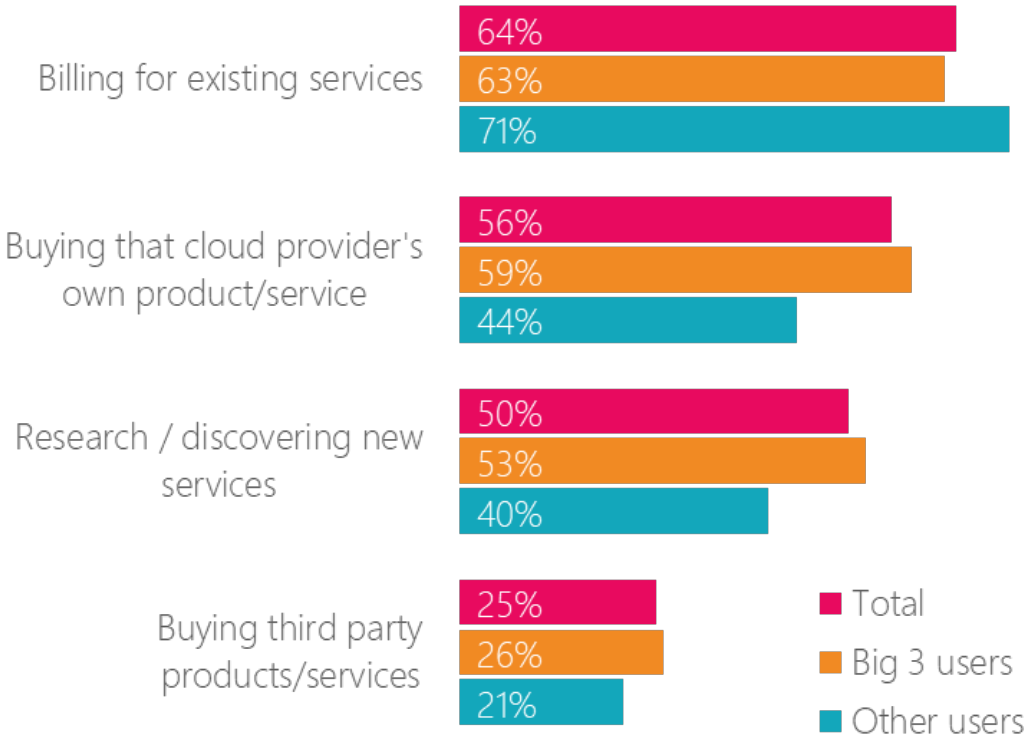


*Source: Context Consulting research data tables, Q46.*

4.15 The most frequently cited reasons for using marketplaces were billing for existing services (64% out of all respondents who use marketplaces), buying the cloud provider's own products (56%), and research and discovering new services (50%).

4.16 Buying third-party products and services was mentioned by 25% of those who use marketplaces (or 13% of all IaaS/PaaS users).[195] Buying third-party products and services is relatively more important for larger companies with 2500+ employees (at 31% of those who use marketplaces), companies older than 20 years (31%), those in public sector (40%), and AWS users (33%).

4.17 Some of the large customers that responded to our statutory information requests told us they were not using marketplaces at all or were using them in very limited cases ([✂]).[196] [✂] told us it very rarely uses marketplaces because its internal processes for authorising such purchases are burdensome.[197] [✂] told us it is their general policy not to purchase cloud-based software from marketplaces as they do not have sufficient cost and contractual controls, meaning it is difficult for [✂] to keep track of transactions and control cost.[198] [✂] noted that there are risks to using marketplaces, as hyperscalers have the ability to remove

[195] Analysis of Context Consulting research data, Q45 & 46.
[196] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[197] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[198] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

a customer's access to marketplace, which could then impact their relationship with third parties. For this reason, [✂] avoid using marketplaces.[199]

4.18 At the same time, some other customers told us they use marketplaces ([✂]),[200] mainly because spend through marketplaces is counted towards securing discounts, and billing and terms are simplified.

## Contracts, pricing, negotiation

4.19 Data from each of AWS and Microsoft suggests that the [✂] of their customers in 2022 were PAYG customers (around [✂]% and [✂]% for AWS and Microsoft, respectively).[201] This is in contrast to our market research findings, which found that the vast majority of respondents had contracts with cloud providers, while only about 17% of user-provider relationships[202] were on a PAYG basis (i.e. respondents). While there may be a range of factors for this difference, it is likely that part of this is down to a different interpretation of the terms. In terms of UK revenue shares, PAYG customers account for [✂]% ([✂])[203] of their revenues.[204]

4.20 We asked the hyperscalers about the length of their spending commitments and agreements with customers. They told us that in 2022:

a) [✂]% of their total UK cloud revenues were accounted for by customers with spending commitments and agreements lasting at least one year;
b) [✂]% of UK revenues were accounted for by customers with spending commitments and agreements lasting at least three years;
c) [✂]% of UK revenues were accounted for by customers with spending commitments and agreements lasting at least five years; and
d) there was more variation in use of agreements longer than five years, which, in terms of UK revenues, accounted for [✂]% of revenues to [✂]% of revenues.[205]

4.21 We asked respondents whether their contract or purchase of IaaS/PaaS cloud services were separate from other IT purchases, or whether these purchases were bundled with other

---

[199] Ofcom / [✂] meeting, [✂].
[200] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to our customer questionnaire, question [✂].
[201] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[202] Respondents were asked the length of the contract for each provider they were using, so if a user had 3 providers, this represents 3 "user-provider relationships" and respondent could, for example, select "5 years" as a response for provider 1, "PAYG (pay-as-you-go)" for provider 2, and "don't know" for provider 3. To summarise these responses, we frame them in terms of "user-provider relationships", or where a contract length was chosen, a contract. Source: Context Consulting research report, slide 94.
[203] We note that for AWS these are UK and Irish revenues ([✂] response dated [✂] to our proposed use of information dated [✂]).
[204] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]. [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂].
[205] One hyperscaler was only able to provide this information for 2021 rather than 2022. [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to follow on questions dated [✂] concerning the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to follow on questions dated [✂] concerning the s.174 notice dated [✂], question [✂]. We note that for AWS these (bullets a-d) are UK and Irish revenues ([✂] response dated [✂] to our proposed use of information dated [✂]).

(non-cloud) products or services. Overall, about four in ten 'user-provider relationships' were reported to have other services bundled in, and in 56% of cases it was a cloud-only purchase.[206] Purchases from AWS were more likely to be stand-alone (62%) than from Microsoft (53%) or Google (52%).

4.22 The most frequently cited reason to buy other services alongside IaaS/PaaS services was cost-effectiveness (74%) followed by assurance that everything works together (51%), and convenience (48%).[207] Five percent of respondents said they had to buy cloud services together with other services.

4.23 Regarding the nature of the pricing used in the relationship with their providers, in just one third of cases (i.e. user-provider relationships) it was on a 'price as quoted' basis, while in 42% of cases respondents negotiated a discount, and in six percent of cases there was a committed minimum spend. In 21% of cases respondents reported receiving a discount because of buying several services from the same provider, and in 10% of cases because of buying some non-cloud services. Our market research found Google's contracts are more often based on price as quoted without any discounts (at 40%). For AWS and Microsoft, the picture is close to the average.

4.24 Findings from our market research show that in one in three cases, IaaS/PaaS customers find it difficult to accurately predict the future costs of their cloud computing.[208] We also heard that 52% of customers were concerned or very concerned about a lack of pricing transparency.[209] Our engagement with customers also indicated varied experience in a customer's ability to predict cloud spend, with some finding it relatively straightforward, and others highlighting challenges.

4.25 In responses to our statutory information requests, some customers told us that there is often very little room for negotiation with the hyperscalers, if at all. [✂] suggested that even large companies like itself do not have a strong negotiating position (e.g. over contract terms and price increases) because of their increasing dependence on single cloud providers.[210] [✂] told us that there are limitations on the extent to which businesses of [✂] size are able to negotiate specific terms with AWS and Microsoft. While there may be room at times for some 'non-standard' terms to be applied, [✂] suggested these are limited and based on spend levels. And while [✂] has some specific agreements in place, these are not bespoke, and when it attempted to amend terms beyond the specific agreement terms available, such approaches have been resisted or rejected out of hand.[211]

4.26 [✂] told us cloud providers are unwilling to accept its standard terms as a contract template, instead requiring it to adopt their standard templates, including policies in relation to data protection and security. [✂] told us the scope of changes is usually limited to price, minimum term and spend commitments, with limited flexibility to negotiate other dimensions, such as service level agreements and technical specifications of their services.[212]

---

[206] Context Consulting research report, slide 89.
[207] Context Consulting research report, slide 89.
[208] Context Consulting research report, slide 95. Findings showed that in about half (52%) of cases, the respondents found it 'quite easy' or 'very easy' to accurately estimate future costs, and in about 30% of cases it was somewhat or very challenging. The results for larger providers are similar to the average.
[209] Context Consulting research report, slide 131.
[210] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[211] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[212] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

4.27 A number of customers ([✂]) say they were able to negotiate bespoke agreements with cloud providers.[213] Usually, discounts depend on contract length and committed spend, and a longer contract and higher spend commitment typically results in a greater discount.

### Use of competitive tenders

4.28 Several cloud providers told us that tenders are often held by clients in the public sector where it can be a regulatory requirement ([✂]).[214] [✂] also suggested it is more common for larger organisations to run competitive tenders.[215] We summarise our understanding of public procurement in Section 3.

4.29 Some customers told us they used tenders before ([✂]), usually for limited specific purposes, while others ([✂]) never used them.[216]

### Contract renegotiation

4.30 In about 58% of user-provider relationships in our market research, the respondent had renewed or renegotiated a contract at some point in the past. For Microsoft Azure, AWS and Google Cloud these numbers are close to the average (55%, 54% and 58%, respectively).

4.31 Most of the customers who responded to our statutory information requests renegotiated their contracts at some point. Usually, it is not the whole contract that is renegotiated, but particular terms related to discounts and committed spend. We also heard from customers and providers that in some instances customers may renegotiate before they reach the end of their contract, for example, if a customer's use of the cloud has grown more quickly than originally anticipated.

## How customers choose between different providers

4.32 The key decision point when choosing a cloud provider is when the customer first migrates into the cloud, which the vast majority of customers we have spoken to agree with. For some customers there can be further decision points, such as when customers phase migration by workload or business unit. As customers' needs develop and evolve, they will need to make further decisions on who provides for incremental use cases. We explore the factors customers consider when making these decisions below.

4.33 Based on the information we gathered in the first half of this study; we identified a number of factors that can determine customers' choice of cloud provider. Following responses to

---

[213] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

[214] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂] question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

[215] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

[216] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

our interim report and stakeholder engagement, we have identified the following factors as a list of considerations that determine customers' choice of cloud providers. We discuss the relative importance of these factors later in this subsection. These factors are not mutually exclusive, and a weighted combination of them will influence a customer's final decision:

a) **Quality and range of services**. In the market research the quality of service is mentioned most frequently as an important factor when choosing a provider, as well as being cited as the most important factor. Overall, quality of service was cited as an important factor when choosing a provider in 39% of cases.[217] Some customers may particularly value 'must-have' services that only certain providers offer, whilst others may look out for the breadth of a providers' service catalogues (number of features was cited as an important factor when choosing a provider in 31% of cases) as they value the convenience of being able to purchase all of their cloud services from a single provider. Furthermore, responses across the board have emphasised the importance of scalability. Customers want the ability to increase usage on demand and have systems respond rapidly and effectively, which can be a key criterion in selecting a provider of cloud services.
b) **Pricing and costs**. The potential to reduce costs by moving to usage-based pricing is commonly cited by customers as a key motivation for using public cloud. The potential to optimise costs is also important as customers continue to increase their spend on cloud.[218] From our market research, 'best value for money' is the second most frequently cited reason for choosing a provider, mentioned in 33% of cases. It seems to be value for money rather than absolute cost that is the more important factor in choosing cloud provider.[219]
c) **Ease of integration**. The ability to easily integrate cloud services with existing IT infrastructure (i.e. traditional IT or private cloud environments) is another relevant factor for customers when choosing cloud providers, as is the time it takes to implement and run a new cloud service. In addition, the ability to run software that interoperates with other cloud services or requires data to be exchanged with another cloud may be important.
d) **Reputation and existing relationship**. Customers must trust that their data and workloads will be secure and accessible. The ability of cloud providers to handle large amounts of data and their track record of service availability may be an important factor for some customers. The level of customer service being offered and established relationships between customers and providers in other markets may also be a consideration. In our market research, supplier reputation is the third most frequently cited factor in choosing a provider, mentioned in 32% of cases.
e) **Geographic reach**. The global reach of a cloud provider, and the availability of local data centres in multiple territories, is important for some customers. The ability to host and process data in certain regions may be important for legal or regulatory reasons. We explore this further below.
f) **Security and resilience**. The security and resilience arrangements of a cloud provider can be a key consideration for customers, particularly where they need to comply with

[217] Context Consulting research report, slide 68.
[218] Flexera, 2023. Flexera 2023 State of the Cloud report [accessed 26 September 2023]. Flexera suggest that respondents anticipate organisational spend on public cloud to grow by 30% in the next year, page 41.
[219] 'Offered the best price' comes only eighth on the list of factors for provider choice in our market research, mentioned in 24% of cases.

relevant regulatory requirements. In our market research, 'proposed level of security' is in the top five most important factors in choosing a provider, chosen in 31% of cases.

g) **Regulatory compliance**. Depending on local data policy and regulation, as well as sector-specific regulation, customers also need to take into account broader obligations and geopolitical factors in choosing a cloud provider. For example, finance or telecoms customers have security and service availability obligations, in addition to broader data processing obligations. Cloud customers will want to be confident that their cloud services provider meets applicable laws and standards, and that they have control of how their data is stored and processed.

h) **Skills**. Availability of skilled resources is also an important factor in provider choice, with our market research respondents considering it one of the key criteria in 22% of cases. The costs associated with retraining or up-skilling existing resource is another important consideration for firms when considering providers, as well as possible organisational impact including challenges with context switching and communication across teams.

4.34 **Other factors**. Considerations such as 'increasing focus on environmental impact' that may lead customers to give weight to the sustainability credentials of different cloud providers. Customers may also be influenced by wider societal and economic changes that may influence their priorities. One stakeholder noted that hyperscalers can offer significant discounts or other benefits, such as skills training, that other suppliers may not be able to match, which, in their view may also impact customer decision making.[220]

## Some customers use a multi-stage decision process

4.35 Several customers told us in response to our statutory information requests that the decision to choose a provider is made in stages:

a) Usually, first there is a set of minimum technical requirements, a provider that does not satisfy these will not be considered further. These requirements may differ from customer to customer.

b) If a provider satisfies the customer's minimum requirements, then customers consider financial factors, including price, discounts, minimum committed spend etc. However, some respondents did not mention this step specifically and were more concerned about technical requirements.

c) If there are still several providers to choose from, other factors come into play, such as additional technical requirements, and customer service.

4.36 Other customers did not describe choosing a provider as a multi-stage process and consider all relevant factors in the round when deciding.

## Relative importance of factors

4.37 Figure 4.4 illustrates the relative importance of factors we described above for respondents to our market research. Customers were asked to select the reasons they chose a specific provider in each case. Whilst it is not clear when interpreting survey results, we think it is most likely respondents are considering the situation when they first chose a provider, as opposed to responding for each additional workload they may have migrated to the cloud. This is an important consideration when interpreting the findings. Across all cases, the top five reasons were service quality (39%), best value for money (33%), supplier reputation (32%), proposed level of security (31%) and number of features (31%). The same reasons were given most frequently when respondents were asked about single most important

[220] Ofcom / [✂] meeting, [✂].

reason, with 'existing relationship for other services' moving up to the third place of most important reasons.

**Figure 4.4: Reasons for choosing a provider, as reported in our market research**

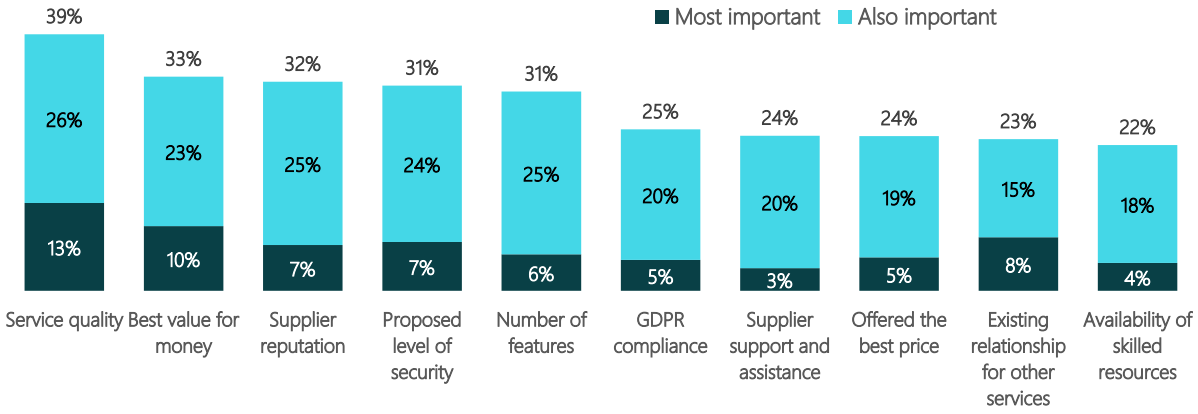


*Source: Context Consulting research report, slide 68.*

4.38 We asked respondents which providers they used and the reasons for using each of those providers. Figure 4.5 shows what factors came into play for different providers. Most factors are common for all providers, with service quality being the most frequently cited reason for many and best value for money and proposed level of security cited regularly for all.

4.39 'Existing relationship for other services' is relatively more important for Microsoft Azure and BT Group customers, while is not in the top six reasons for other major providers. 'Supplier support' is among the top factors for Google, IBM and BT. Only for Oracle customers did 'availability of skilled resources' reach the top six reasons.

**Figure 4.5: Top reasons for choosing a provider, by provider, as reported in our market research**

| | Microsoft Azure | AWS | Google cloud | Oracle cloud | IBM cloud | BT |
|---|---|---|---|---|---|---|
| 1 | Service quality | Service quality | Service quality | Service quality | Service quality | Best value for money |
| 2 | Supplier reputation | Best value for money | Best value for money | Proposed level of security | Proposed level of security | Service quality |
| 3 | Proposed level of security | Proposed level of security | Number of features | Supplier reputation | Supplier reputation | Existing relationship for other services |
| 4 | Best value for money | Number of features | Supplier reputation | Availability of skilled resources | Number of features | Supplier reputation |
| 5 | Number of features | Supplier reputation | Proposed level of security | Best value for money | Best value for money | Proposed level of security |
| 6 | Existing relationship for other services | Offered the best price | Supplier support and assistance | Number of features | Supplier support and assistance | Supplier support and assistance |

*Source: Context Consulting research report, slide 70.*

4.40 We also asked customers (for each provider they used) whether, at the time when they selected that specific provider, they considered other options. Overall, in 52% of the cases the respondents said they had considered other providers and they had a range of options; in 36% of the cases they did consider other providers but their options were limited; and in

9% of the cases they had only one feasible option. Users of AWS, Oracle and IBM were more likely to say they had a full range of options (at 57%, 55% and 58%, respectively) compared to the average customer experience. More broadly we asked customers how much competition they thought there was in IaaS and PaaS. Overall, the slight majority of users felt there was some competition in IaaS but not so in PaaS.[221] While customer's perceptions of competition are interesting, given the different ways in which such questions can be interpreted we place more emphasis on our analysis of more objective measures of market outcomes set out below.

## How customers choose between deployment models

4.41 Above, we presented the factors that are considered important to customers when choosing between cloud providers. As well as potentially determining a customer's choice of cloud provider, some of these factors (i.e. security and resilience) may also be important in determining a customer's choice of deployment model. In Section 3, we outlined evidence on the customer use of different cloud deployment models (public cloud, private cloud, hybrid cloud), and we also recognised that some customers may use traditional IT (on-premises) as another deployment model.

4.42 We recognise that there may be some substitutability between public cloud and the other deployment models. In particular, some cloud providers (e.g. [✂]) indicated that they do not distinguish their products across the different cloud deployments that they cater to (public, private and hybrid) and that customers have the ability to deploy relevant products in any of these models.[222] Also for certain use cases, such as storing data, some cloud providers (e.g. [✂]) suggested that customers can use cloud and on-premises IT solutions to run the same workloads and address the same requirements.[223]

4.43 In their response to the interim report, both AWS and Microsoft presented the cloud market within the context of a wider IT services spend.[224] [✂] said that a customer is typically looking to solve a specific IT problem and will choose among many options, whether it be on premises, in a co-located environment, online or adopting a hybrid approach using multiple of these options.[225] [✂] told us that "from a customer's perspective, IT services are substitutable for most use-cases, regardless of their delivery method".[226] Furthermore, [✂] said that they believe private cloud and on-premises traditional IT solutions are credible substitutes for public cloud.[227]

4.44 In its submission dated [✂], [✂] pointed to a survey conducted by Public First and told us that the results from the survey on respondents' switching behaviour suggest that cloud

[221] Context Consulting research report, slide 132.
[222] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[223] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[224] AWS response to the interim report pages 2-3; Microsoft response to the interim report page 7. Microsoft also cited a recent Amazon earnings call [accessed 16 August 2023].
[225] [✂].
[226] [✂].
[227] [✂].

service providers face significant competitive pressure from on-premises providers.[228] As explained in more detail in Annex 3, we disagree with [✂] argument, because (amongst other reasons) we think that: this historical evidence is not suitable to shed light on substitutability between on-premises solutions and cloud services at the margin; and it is not clear whether this is capturing switching from public cloud to on-premises, especially given the relatively high proportion of IaaS/PaaS users who reported using private cloud only.

4.45 We recognise that customers looking for a new product would consider the most appropriate deployment model to facilitate that product. However, we think the extent of substitutability is likely to be limited due to the distinct characteristics of each deployment model. Customers, as well as cloud providers themselves, recognise that the resulting choice of deployment model often depends on the specific needs and requirements of individual workloads of a customer.

4.46 Generally, once a customer has decided to migrate a workload into the public cloud, we think alternative deployment models are likely to be a weak substitute for these workloads, in most cases. Whilst it is possible for a customer to use an on-premises solution or a public cloud solution, we have identified a number of challenges associated with using on-premises solutions that make it only a weak substitute in practice. [✂] noted the lengthy procurement cycles involved in an on-premises model compared to accelerated digital product development available in cloud.[229] AWS explained that customers are often 'locked-in' to expensive on-premises infrastructure and that switching can be prohibitively expensive and time consuming.[230] This highlights cost as one of the key reasons behind the move to cloud for many customers, and why on-premises solutions are not always a close, or even viable, substitute.

4.47 Moreover, we think that the extent of substitutability is likely to be limited due to the distinct characteristics of each deployment model. We observe that customers:

a) Use public cloud to often meet requirements such as scalability and/or innovation (e.g. Sainsbury's, ITV, Netflix).[231] These advantages are a key reason why customers are increasingly migrating workloads from on-premises IT to public cloud. Both Sainsbury's and ITV suggested that the use of public cloud gives them the flexibility to respond to customer needs and scale up in times of increased demand.[232] In particular, Sainsbury's noted that they would not have been able to respond to the rapid changes in customers' online demands, during the pandemic, if their critical business systems had not been in the public cloud.[233] Similarly, Netflix suggest that the elasticity of cloud has supported their rapid growth over time. They also note that since moving to public cloud, their

---

[228] [✂]. The results of the Public First's survey are available at: CCIA survey [accessed 7 September 2023]. In the Public First's survey, only a minority of the 716 IaaS/PaaS users included in the survey – 26%, corresponding to 185 out of 716 IaaS/PaaS users – had switched cloud provider in the past few years. Within this group of "switchers", the vast majority – 69%, i.e. 127 out of the 185 "switchers" – had switched to another cloud provider, whereas 29% – i.e. 54 out of the 185 "switchers" – had moved to an on-premises solution].

[229] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].

[230] AWS response to the interim report, page 4, paragraph 9.

[231] Sainsbury's response dated 8 December 2022 to our customer questionnaire, question 19; ITV response dated 15 December 2022 to the s.174 notice dated 24 November 2022, question 1; and Netflix, 2016. Completing the Netflix Cloud Migration [accessed 25 September 2023].

[232] Sainsbury's response dated 8 December 2022 to our customer questionnaire, question 19; and ITV response dated 15 December 2022 to the s.174 notice dated 24 November 2022, question 1.

[233] Sainsbury's response dated 8 December 2022 to our customer questionnaire, question 19.

"costs per streaming start" has reduced to a fraction of the costs previously incurred when relying on their in-house data centres.[234] FirstGroup plc migrated workloads to AWS from an on-premises solution to "improve performance and reliability". [235]

b) Use private cloud to meet specific requirements around latency, security, resilience and/or regulatory compliance (e.g. [✂][236]). For example, all mobile network operators in the UK currently use private cloud to run their major network workloads, due to security and resilience requirements. We note that private cloud may generally be attractive to industries that require high security and resilience (e.g. telecoms, financial services). Private cloud is also used by some broadcasters, for example when managing video and audio streams for live productions, for which low latency may be particularly important.

c) Use hybrid cloud where there is a need for a mix of both public and private cloud to support different use cases, and/or to support legacy applications that cannot easily be re-architected to work in the public cloud (e.g. [✂][237]); and

d) Use on-premises solutions where there is a specific customer need to maintain control over physical hardware – for example this could be for commercial reasons, security or regulatory requirements (e.g. [✂][238]). In 2016, Dropbox decided to move out of the cloud and build its own data centres, due to a need for more control over the underlying hardware and infrastructure.[239]

4.48 Furthermore, CMA research showed that customers it spoke to "considered that the benefits of moving certain existing workloads from enterprise deployments to the public cloud typically outweighed the costs" with the most common driver being "the cost effectiveness of public cloud relative to traditional data centres".[240] Therefore, we expect for many customers, once they have migrated workloads to public cloud, they are unlikely to consider switching them back to on-premises. The inherent benefits of on-premises solutions are slowly reducing over time as innovation and scalability in the cloud makes it a distinctly different product to other deployment methods.

4.49 There are instances where customers may switch between these deployment models, dependent on the extent to which a given application is designed to take advantage of the unique environment provided by that deployment model. Stakeholders provided some limited examples of customers switching between on-premises and the cloud, including Dropbox who moved a lot of their data out of AWS to their own network of servers, and Zynga a mobile game developer, who switched from AWS to their own private cloud and back again.[241] Based on the sum of evidence we have gathered, we believe these examples are quite specific to the circumstances of those customers and are the exception rather than

---

[234] Netflix, 2016. Completing the Netflix Cloud Migration [accessed 25 September 2023].
[235] AWS response to the interim report, paragraph 12
[236] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[237] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[238] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[239] TechCrunch, 2019. DropBox infrastructure continues to evolve [accessed 25 September 2023].
[240] CMA, 2023. Anticipated acquisition by Broadcom Inc. of VMware, Inc. Final Report. page 58, paragraph 7.60 [accessed 25 September 2023].
[241] AWS response to the interim report, paragraph 12.

the rule. For example, Dropbox moved due to public cloud not providing enough control over underlying hardware, and Zynga switched back because it found that keeping its own infrastructure was costlier over time than using AWS. [242]

4.50 Overall, given the evidence we have gathered during this market study, we think that there is limited substitutability between public cloud and alternative deployment models. Accordingly, we consider it appropriate to focus our competitive assessment in remaining sections on providers of public cloud (unless otherwise stated).

# Prevalence of multi-clouding and switching by customers

## Prevalence of multi-clouding

4.51 In their responses to our interim report, both AWS and Microsoft referred to the Context Consulting market research and other public reports (discussed in Annex 3) as evidence of the prevalence of multi-cloud.

4.52 The market research we commissioned from Context Consulting indicated that 52% of respondents reported using more than one IaaS/PaaS provider. However, as noted in our interim report, this share is likely to overstate the actual proportion of customers that use more than one public cloud provider. This is because, in addition to using multiple public cloud providers, respondents who use more than one IaaS/PaaS provider might have included in their response:

a) the products of an ISV and public cloud provider on the same cloud;
b) private and public cloud solutions (i.e. hybrid cloud); or
c) two private cloud providers.[243]

4.53 Other public surveys – such as the ones carried out by Flexera, Oracle, Foundry and Public First – report higher shares of respondents using multi-cloud, although several reasons may explain the differences with the Context Consulting research (e.g. the surveys asked about SaaS in addition to IaaS/PaaS or the respondents were mostly large enterprises).[244]

4.54 Our market research also asked respondents who use more than one IaaS/PaaS provider about how they multi-cloud. According to Context Consulting, 45% of this subset of respondents use different IaaS/PaaS providers for different workloads; 40% spread similar workloads across IaaS/PaaS providers; and 15% have one main IaaS/PaaS provider and use others as a back-up.[245] We remain cautious about mapping these shares on to the three types of multi-cloud architecture we use in this report. The 45% of respondents who use different IaaS/PaaS providers for different workloads might be using siloed multi-cloud; and the 15% of respondents have one main IaaS/PaaS provider and use others as a back-up might be using cloud duplication. However, the 40% of respondents who spread similar workloads across IaaS/PaaS providers might be using siloed and/or integrated multi-cloud. We are also of the view that, in general, respondents might have overstated their use of multiple cloud providers (as noted above).

4.55 Among those customers who responded to our requests for information during the market study, use of integrated multi-cloud was very limited. [✂] told us that additional complexity

[242] AWS response to the interim report, paragraph 12.
[243] Context Consulting research report, slide 75. See also Annex 3 for additional details.
[244] See Annex 3 for references and additional details.
[245] Context Consulting research report, slide 79. See also Annex 3 for additional details.

and reduced optimisation mean that integrated multi-cloud is only used in a very limited number of cases.[246] Another customer, [✂], told us that they duplicate their databases on AWS and Google Cloud, although, for cost reasons, the back-up cloud of Google is more basic (albeit capable of being scaled up at speed, if necessary).[247] Other customers – for example, [✂],[248] [✂] [249] and [✂] [250] – told us that they were using different cloud providers for separate workloads (e.g. Google for specialised use cases, such as big data and analytics).

4.56 In response to our interim report, [✂] submitted a quantitative analysis of their UK customers' cloud usage and opportunity data. According to [✂], the result of this analysis show that, on a conservative basis, approximately [✂]% of [✂] UK customers (weighted by revenue) in the data awarded a tender to at least one other cloud provider between 2017 and 2022, indicating that they were multi-clouding.[251] However, as we explain in detail in Annex 3, we consider that the approach used by [✂] does not provide a meaningful measure of the prevalence of multi-cloud, as it only takes into account [✂] customer spend on its own services. Moreover, [✂] approach gives more weight to its largest customers. When we do not weigh customers by revenue, we find that only [✂]% of [✂] UK customers use another cloud provider in addition to [✂] – a proportion which is significantly lower than the one reported by our Context Consulting market research. Lastly, we have reservations about the quality and accuracy of the data used by [✂] in its analysis – which [✂] also acknowledges.[252]

4.57 Overall, while many customers use more than one IaaS/PaaS provider, we remain of the view that integrated multi-cloud is uncommon.

### Prevalence of switching

4.58 Our Context Consulting market research also included questions about the prevalence of switching. Among IaaS/PaaS users, Context Consulting found that 18% of respondents had switched IaaS/PaaS providers completely (and stopped using the previous IaaS/PaaS provider) in the past; 35% had taken on an additional IaaS/PaaS provider; 35% had considered switching but did not switch; and 23% had never considered switching.[253] We are cautious about these findings, since – as we explain in Annex 3 – they may overstate the actual level of switching taking place in the market.

4.59 In its response to our interim report, [✂] submitted a quantitative analysis based on its UK customer usage and revenue data, which would suggest, according to [✂], that its customers can and do switch cloud provider.[254] However, as we explain in detail in Annex 3, the churn rates calculated by [✂] – [✂]% across all cloud services and [✂]% considering cloud compute services only – are similar to those of the Context Consulting market research and consistent with a finding that switching levels are low in the cloud market.

4.60 Furthermore, we have reservations about the approach used by [✂]. Firstly, the churn rate calculated by [✂] is based on reductions in spending by customers (as opposed to full

---

[246] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[247] [✂] response dated [✂] to our customer questionnaire and Ofcom / [✂] meeting, [✂].
[248] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].
[249] [✂] response dated [✂] to the s.174 notice dated [✂] question [✂].
[250] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].
[251] [✂]; and [✂].
[252] See Annex 3 for additional details.
[253] Context Consulting research report, slide 105. See also Annex 3 for additional details.
[254] [✂]; and [✂].

switching) and it may only imperfectly capture switching – something that [✂] also acknowledges. Moreover, once smaller customers – e.g. those spending less than $[✂] or $[✂] in total during the entire period under consideration (and so, in our view, may have just tried [✂] services or only used those services occasionally) – are removed from the analysis, the churn rates for [✂] customers (both across all services and cloud compute services only) become significantly smaller.[255]

4.61 Overall, we remain of the view that switching levels are low in the cloud market.

# How providers compete

4.62 In this subsection, we consider the extent to which customer preferences are reflected in how suppliers compete to attract and acquire customers. The cloud services supply chain is complex, with different types of suppliers providing services at some or all levels of the cloud stack. We assess the key dimensions of competition between cloud providers: pricing, and quality (in terms of range of services, innovation and ease of integration).

4.63 We first outline how providers compete through pricing strategies, and on other measures such as quality and range of services. We then set out the sales channels that some cloud providers and ISVs use, before discussing the role of ecosystems in how providers compete in cloud.

## Pricing and costs

4.64 The potential to reduce costs is commonly cited by customers as a key reason for migrating towards public cloud. However, the initial process of migrating workloads to the cloud can be particularly costly for some customers. Cloud providers compete to attract new customers by providing a range of pricing benefits in the form of discounts (including committed spend discounts), cloud credits and free trials/tiers, which contribute towards lowering the initial migration costs that customers may face.[256]

4.65 **Discounts** are typically dependent on factors such as type of usage, contract length and/or minimum spend. Committed spend discounts (i.e. where customers commit to consuming a minimum monetary value of services over an agreed period) can either apply to individual products or families of products (the committed spend discounts available to all customers are generally of this type), or to a customer's total spend with a cloud provider (these tend to only be privately negotiated). Microsoft and AWS state such commitments provide them with a degree of certainty in relation to capacity requirements, contributing to costs savings, which can then be passed on to customers in the form of discounts.[257] Cloud providers also offer other types of discounts. For example, discounts may be offered, through the adoption of pricing models such as 'spot instances', which enable customers to take advantage of spare capacity at discounted rates relevant to standard, PAYG prices.[258]

[255] See Annex 3 for additional details.
[256] We note that some of these incentives (e.g. discounts) are not exclusively provided to customers migrating to the cloud. They may also be provided to existing cloud customers when renewing contracts, or to attract customers from other cloud providers.
[257] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and AWS response to the interim report, page 13, paragraph 31.
[258] AWS website. Amazon EC2 Spot Instances [accessed 26 September 2023]; and Microsoft website. Azure Spot Virtual Machines [accessed 26 September 2023].

4.66 **Cloud credits** provide customers with a spending allowance on eligible cloud services. Most cloud providers offer credits to customers when they open an account with them for the first time. These credits are typically around $200 and are valid for one or two months.[259] [✂] and [✂] suggest they provide credits so that new customers can test and explore their cloud services.[260] New customers may also be eligible for higher credit amounts as part of credit programs offered to specific groups of customers. For example, AWS offers up to $100k, Microsoft offers up to $150k, and Google offers up to $100k for each year over two years (so a total of $200k), as part of their credit programs for 'start-ups'.[261] We are also aware that cloud credits may be offered to customers on a case-by-case basis and therefore there may be scope for customers to negotiate and benefit from higher credit amounts.[262]

4.67 **Free trials/tiers** provide customers with services that are free of charge for a specified usage and limited period of time. For example, AWS provides 2,200 instance hours per month of 'Amazon ECS Anywhere' for free over a period of 6 months.[263] Free trials/tiers tend to be available to customers when using a cloud provider's services for the first time, with customers typically viewing them as an inducement to try their services.[264]

## Quality and range of services

4.68 Customers also consider quality and range of cloud services as important factors when choosing a specific cloud provider. This is also reflected in how cloud providers, in particular the hyperscalers, compete with each other by increasing their range and functionality of cloud services.

4.69 **Range of services** – the hyperscalers tend to view their ability to provide a range of services as a defining feature of their strength in attracting customers, and in competing with other cloud providers.[265] One stakeholder said that they changed business model from providing cloud services to public sector organisations to exclusively focus on managing more sensitive data. It said this was in response to hyperscalers entering and rapidly winning significant numbers of public sector contracts in the stakeholder's usual area of operation.[266]

4.70 **Quality and innovation –** cloud providers consider quality and innovation as important to remaining competitive, and in meeting the needs of customers in cloud services. For example, Capgemini observe that the hyperscalers maintain "a very high rate of innovation".[267] By continuing to innovate, cloud providers can improve their performance

---

[259] Microsoft website. Azure free account [accessed 26 September 2023]; and IBM website. IBM Cloud free tier [accessed 26 September 2023].
[260] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[261] AWS website. AWS activate [accessed 26 September 2023]; Microsoft website. Unlocking Azure credits as your start-up grows [accessed 26 September 2023]; Google website. Google for start-ups cloud program [accessed 26 September 2023].
[262] Such credit offerings may also be linked to spend commitments. [✂] response dated [✂] to the s.174 notice [✂], question [✂].
[263] AWS website. AWS Free Tier [accessed 26 September 2023]. We note that cloud providers may also additionally provide certain cloud services as free of charge always (permanent free tier).
[264] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[265] Ofcom analysis of [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[266] Ofcom / [✂] meeting, [✂].
[267] Ofcom / Capgemini meeting, 1 November 2022.

and technical capabilities, therefore enabling them to provide a better quality of service for customers.

4.71 **Ease of integration –** interoperability facilitates building solutions which require integration of several products, so cloud providers seek to make integration between their first-party products as simple as possible. Our research indicates that customers care about the level of interoperability offered by cloud providers.[268] We think interoperability is important because it impacts a customer's ability to switch and/or build their preferred cloud architecture by combining services from multiple clouds.

## Geographic coverage

4.72 In our market research we found that geographic reach of infrastructure was chosen by a smaller share of respondents as important when choosing a cloud provider in comparison to other factors.[269] It may be less of a differentiating factor between providers because, as detailed later in this section, several providers now have a similar global reach. However, geographic coverage will be more important for some customers that have specific needs and requirements, such as:

a) **Performance reasons:** the latency of cloud services is an important consideration for some customers, e.g. where the customer provides services to end-users that require high availability and fast response times such as video streaming or a payments app. To reduce latency, customers will consider whether cloud providers can host and run their workloads close to their end-users. The geographic coverage of infrastructure can also influence the resilience of cloud services.[270]

b) **Data sovereignty:** some customers may face legal or regulatory restrictions on where they store and transfer their data. In addition, certain customers may have internal policies where there is a preference to store data locally. Hence, the location of datacentres may be a relevant factor for some customers when choosing providers.

c) **Ease of procurement for global companies:** multi-national companies operating in a number of regions may prefer to use a single provider that covers all of the regions they operate in. Using a single provider may offer some advantages such as simpler procurement and pooling of internal skillsets.[271] Companies operating in different regions across the world may also look for a provider that has datacentres located in close proximity to their customers, in order to provide services with sufficiently low latency, or to minimise data transfer costs (as discussed in Section 5).

---

[268] Context Consulting research report, slide 131. 52% of respondents were concerned or very concerned about the lack of interoperability.

[269] 'Location of the data centres of the supplier' was cited as an important reason for choosing provider in 14% of cases. In comparison, the top answer – 'service quality' – was cited as an important factor in 39% of cases. However, we consider that global reach can also impact service quality, e.g. resilience and latency, therefore the importance of global reach may also be reflected in the importance of service quality. Context Consulting research report, slide 69.

[270] For example, in our qualitative research we heard from a customer that they "insist on two geographically diverse sites with at least dual internet going into it, dual power supply going into each etc." Context Consulting research report, slide 31.

[271] In our qualitative research, one respondent considered that only Google, AWS and Microsoft have the reach to meet the needs of their global business. Context Consulting research report, slide 64.

## Sales channels

4.73 Broadly consistent with our section above on how customers purchase cloud services, the majority of cloud services are bought directly from providers, with some bought from third parties and a small proportion from marketplaces. Table 4.6 sets out the proportion of sales by sales route for some cloud providers and ISVs in 2022. The importance of different routes varies, but for most providers the direct sales route is the largest. Some providers make greater use of professional services providers (particularly [✂]). Sales through first- and third-party marketplaces are in general a small proportion of providers' total sales revenue.

**Table 4.6: Estimated proportion of UK sales (by revenue) through different sales channels, 2022**

| | Direct sales | Suppliers of professional services | Marketplaces |
|---|---|---|---|
| **AWS** | [✂]% | [✂]% | [✂]% |
| **Google** | [✂]% | [✂]% | [✂]% |
| **Microsoft** | [✂]% | [✂]% | [✂]% |
| **IBM** | [✂]% | [✂]% | [✂]% |
| **Oracle** | [✂]% | [✂]% | [✂]% |
| **OVHcloud** | [✂]% | [✂]% | [✂]% |
| **MongoDB** | [✂]% | [✂]% | [✂]% |
| **Snowflake** | [✂]% | [✂]% | [✂]% |
| **VMware** | [✂]% | [✂]% | [✂]% |

*Sources: AWS response dated 31 March 2023 to s.174 notice of 24 October 2022, Part B Q5, and Annex 2; Google response dated 31 March 2023 to s.174 notice of 26 October 2022, Part B Q5, Annex 2; IBM response dated 4 April 2023 to s.174 notice of 25 October 2022, Part B Q5, Annex 2; Microsoft response dated 25 April 2023 to s.174 notice of 21 October 2022, Part B Q5, Annex 2; MongoDB response dated 31 March 2023 to s.174 request of 27 October 2022, Part B Q5, Annex: Part B Cloud Services Template updated; Oracle response dated 31 March 2023 to s.174 notice of 31 October 2022, Part B Q5, Annex 2; OVHcloud response dated 28 March 2023 to s.174 notice of 27 October 2022, Part B Q5, Annex B; Snowflake response dated 11 April 2023 to s.174 request of 26 October 2022, Part B Q5, Annex 2; VMware response dated 30 March 2023 to s.174 notice of 27 October 2022, Part B Q5 and Annex. Notes: [✂].*

## Ecosystems in cloud services

4.74 In general, ecosystems can be understood as a collection of complementary products and services that work together to create utility for customers. These also typically include an interface or gateway that acts as an intermediary to other components of the market, such as customers, hardware producers and software developers. We believe some components of ecosystems are evident in how some cloud providers operate within the cloud market. Therefore, ecosystems are a relevant framework through which to examine competition within this market. There are several cloud providers that offer ecosystems, but our focus here is on the hyperscalers as the main providers in the UK.

4.75 Each of the hyperscalers offers a wide range of services across all levels of the cloud stack, including both first-party and third-party services (services developed by the hyperscalers themselves and by others, e.g. ISVs). Customers buy a solution of complementary products that work together, meaning purchase decisions for one product can have an impact on purchase decisions of others in the ecosystem. The hyperscalers operate unique cloud environments, collections of programming languages, application frameworks and APIs that allow services across the stack to work together with others within the ecosystem. Hyperscalers also operate marketplaces to allow customers to identify and purchase first and third-party services to match their needs.

4.76 In response to our interim report [✂] told us that they did not agree with the use of the term 'ecosystem' to describe their product offering as they "[do] not prohibit customers in any way from purchasing services from different providers".[272] [✂] told us that they "[strive] to build a high-quality service so that customers *want* to stay with [them]" rather than building a system that locks customers into an ecosystem. Microsoft also told us that the cloud market is not "a digital market characterised by significant network effects or ecosystem lock-in".[273] Microsoft also stated that a customer's ability to switch cloud infrastructure service provider, as well as multi-cloud means that the term 'ecosystem' is inappropriate and that there are no direct network effects in the cloud market.[274]

4.77 In our view, ecosystems are not defined by whether or not customers have the freedom to buy from other providers. Rather ecosystems provide customers with access to a broad portfolio of their own and others' products in a single place that work together seamlessly. We think this characterisation accurately reflects the product portfolios of the hyperscalers, and therefore remain of the view that ecosystems remain a relevant framework to assess competition in cloud infrastructure services.

### Hyperscalers offer a wide portfolio of first and third-party services across the cloud stack

4.78 As discussed above, cloud providers compete on the range of services they offer. Hyperscalers are present throughout the entire cloud supply chain. They offer a strong portfolio of services across the stack, with our analysis showing AWS and Microsoft accounting for around [✂]% [80% to 90%] of UK IaaS revenues in 2022 and [✂]% [50% to 60%] of UK PaaS revenues.[275] Whilst there is more diversity in the SaaS market the hyperscalers still provide a wide range of first and third-party services in this area.[276]

4.79 Customers have told us they value the range of products offered by the hyperscalers,[277] and this is reflected in customer purchase data, with the average hyperscaler customer

---

[272] [✂].
[273] Microsoft response to the interim report, page 37, paragraph 139
[274] Microsoft response to the interim report, page 37, paragraph 139
[275] Ofcom analysis of data provided in response to our statutory information requests and data from Synergy and IDC. This is set out in more detail in the below subsection on UK shares of supply.
[276] Ofcom analysis of data from IDC. This is set out in more detail in the below subsection on UK shares of supply.
[277] [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; and [✂] response dated [✂] to the s.174 notice dated of [✂], [✂].

purchasing multiple services.[278] The range of services that the hyperscalers offer make it easier to buy multiple products from the same provider. This can bring some benefits to customers such as the availability of 'off-the-shelf' services which seamlessly integrate with their existing resources,[279] streamlined procurement processes,[280] and can simplify customers' recruitment and training processes as their staff only need to specialise in a single cloud ecosystem.[281]

4.80 The hyperscalers respond to customer demand for a broad range of services by marketing themselves as the only place customers need to go at any stage of their cloud journey, by being a place for customers to learn about a product, evaluate its appropriateness for their needs, and access ongoing product support.[282] Hyperscalers' websites provide a clear demonstration of this in practice, with an extensive range of information, guidance, and support being offered by each.

4.81 The hyperscalers offer a wide range of first-party products and services, including covering all PaaS segments. They also offer a range of products and services of third-party ISVs which are built on top of the hyperscaler's unique cloud environment, including the collection of programming languages, application framework and set of APIs that allow services across the stack to work together with others within the ecosystem. Customers can therefore buy a solution of complementary products that work together. This means purchase decisions for one product can impact on purchase decisions of others in a cloud provider's portfolio.

4.82 Offering ISV solutions on their infrastructure allows the hyperscalers to significantly broaden the range of services offered within their portfolio, and ensure they are able to meet the needs of a greater number, and variety, of customers. This is particularly the case where ISV products act as complements to the hyperscalers' own product range. However, this relationship between the hyperscalers and ISVs can be complex, because some ISVs offer products which act as direct competitors to the hyperscalers in PaaS and SaaS. We explore the implications of this in more detail in Section 7.

4.83 For the purpose of this market study, we refer to the combination of first and third-party products available on a provider's cloud as ecosystems.

## Hyperscalers can act as distributors of ISV services

*Co-selling*

4.84 Each of the hyperscalers operate co-sell schemes.[283] Co-selling schemes can include a variety of features such as creating visibility for the ISVs solutions with the hyperscalers' sales teams

---

[278] [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[279] [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; and [✂] response dated [✂] to our customer questionnaire, question [✂].
[280] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[281] [✂] confidential response to the CFI, page [✂]; and [✂] response dated [✂] to our customer questionnaire, question [✂].
[282] AWS website. Chapter 3: Customer Engagement Solutions [accessed 1 September 2023]; Microsoft website. What is Azure Marketplace? [accessed 1 September 2023]; Google Website. What is Google Cloud Marketplace? [accessed 1 September 2023].
[283] AWS website. AWS ISV Accelerate Program (amazon.com), and AWS ISV Accelerate Helps Partners Co-Sell with AWS and Reach New Customers | AWS Partner Network (APN) Blog (amazon.com) [accessed 26

and incentivising those sales teams to sell the ISV solutions. This can also extend to greater visibility for PaaS ISVs by making third-party software available alongside the hyperscaler's first-party PaaS. These schemes may also incorporate benefits to the ISV to incentivise them to sell solutions through the provider's marketplace, such as reduced commission fees or including spend on ISV solutions via the marketplace within customer cloud spend commitments.[284] This can include apps which already combine products from ISVs and the hyperscalers, as seen with AWS's partnership with Salesforce.[285]

4.85 The hyperscalers often tailor the level of support provided to ISVs to reflect the level of integration between the ISV and the cloud provider.[286] Access to Microsoft tiers, which offer higher levels of support, requires tight technical integration with Azure and driving Azure-based sales revenue.[287] We also understand that AWS has [✂].[288]

*Marketplaces*

4.86 Each hyperscaler operates a cloud marketplace. A cloud marketplace is an online platform allowing cloud providers and ISVs to sell their services to business customers. Table 4.7 below sets out key data in relation to the UK marketplaces of each of AWS, Google and Microsoft.

4.87 [✂] operates the largest marketplace [✂]. ISVs transact significantly more ([✂]) on [✂] than on [✂]. In addition, [✂] and is [✂].[289]

4.88 The bulk of services offered through marketplaces are those of third parties. AWS only offers a very small number of its own services on its marketplace – and almost all ([✂]%) gross sales revenue on its marketplace is from third-party sales, with listing fee revenue from third-party sales accounting for [✂]% of its marketplace revenues.[290] Our market research indicated that AWS users are more likely to use marketplaces to buy third-party services

---

September 2023]. Google website. Google Cloud doubles-down on ecosystem in 2022 to meet customer demand | Google Cloud Blog [accessed 26 September 2023]. Microsoft website. Sell with Microsoft, Co-sell with Microsoft sales teams and partners overview - Partner Center | Microsoft Learn, and Microsoft Business Applications Independent Software Vendor (ISV) Connect Program onboarding guide | Microsoft Learn [accessed 26 September 2023].

[284] AWS website. AWS ISV Accelerate Program (amazon.com) [accessed 26 September 2023]. Microsoft website. Marketplace rewards - your commercial marketplace benefits - Marketplace publisher | Microsoft Learn [accessed 26 September 2023]. Google website. Google Cloud doubles-down on ecosystem in 2022 to meet customer demand | Google Cloud Blog [accessed 26 September 2023].

[285] AWS website. AWS and Salesforce - Global Strategic Partnership [accessed 26 September 2023].

[286] [✂] explains that it wants to "be the easiest, most efficient, and gainful (new customers) partner to co-sell with our most invested" (meaning those who are most leaned in on [✂]) ISVs ([✂] response dated [✂] to the s.174 notice dated [✂], [✂]). [✂] notes that "The intent [in 2022] is to have a consistent and fair set of partner benefits and compensation treatment for all ISVs [✂] in the same tier to prioritize those partners that will have the highest business impact for customers" ([✂] response dated [✂] to the s.174 notice dated [✂], [✂]).

[287] To achieve Azure IP co-sell incentive status, a solution must reach the required revenue threshold of $100,000 of Azure Consumed Revenue over the past 12 months, and pass the Microsoft technical validation for an Azure-based solution which must confirm more than 50% of an offer's infrastructure uses repeatable IP code on Azure. See Microsoft website. Co-Sell requirements [accessed 26 September 2023].

[288] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].

[289] [✂] response dated [✂] to the s.174 notice dated [✂]; [✂]; and [✂] response dated [✂] to the s.174 request dated [✂], [✂].

[290] Ofcom analysis of AWS response dated 9 December 2022 to the s.174 request dated 24 October 2022, question 19 and Annex 2. Further, AWS was only able to provide a combined figure for first-party services and second-party services (services which AWS purchases and then resells via its marketplace) – and noted that first-party services would be a small subset of this combined figure.

than all marketplace users.[291] Microsoft offers a larger selection of its own services on its marketplace than AWS does on its marketplace. While Microsoft was unable to provide its revenue from first-party sales through Azure Marketplace, our market research indicated that Microsoft users were more likely to use marketplaces to buy first-party services than all marketplace users.[292]

**Table 4.7: Gross sales revenue and commission revenue received from third-party products through AWS, Azure and Google Marketplaces, in the UK (2022)**

| | Gross sales revenue of first-party products | Gross sales revenue of third-party products | Commission revenue from sales of third-party products |
|---|---|---|---|
| **AWS Marketplace** | £[✂] | £[✂] | £[✂] |
| **Azure Marketplace** | £[✂] | £[✂] | £[✂] |
| **Google Marketplace** | £[✂] | £[✂] | £[✂] |

*Sources: Figures for AWS: [✂]. Figures for Microsoft: [✂]. For Google: Gross sales of third-party services data from [✂]. Figures converted to GBP using average USD exchange rate for 2022. AWS was only able to provide a combined figure for first-party services and second-party services (services which AWS purchases and then resells via its marketplace) – and noted that first-party services would be a small subset of this combined figure.[293] Google and Microsoft were [✂].[294] Gross sales refers to the total value of the sales transacted through the marketplace. For Google, these gross sales are net of credits.*

4.89 Overall, we estimate that less than 10% of total ISV sales (PaaS and SaaS) in the UK were transacted via a hyperscaler marketplace in 2022.[295] At present, marketplaces are not a major revenue source for the hyperscalers, accounting for a very small proportion of their total cloud revenues (c.[✂]%).[296] However, there are some indicators that marketplaces may grow in importance over time. Some commentators expect that marketplaces will

[291] Context Consulting research data tables, Q46.
[292] Microsoft response dated 9 December 2022 to the s.174 request dated 21 October 2022, question 19; Microsoft response dated 16 January 2022 to our follow-up email dated 20 December 2022 concerning the s.174 request dated 21 October 2022; and Context Consulting research data tables, Q46.
[293] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to s.174 notice dated [✂], question [✂].
[294] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], page [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], page [✂].
[295] We calculated two estimates of total PaaS + SaaS market size excluding the hyperscalers. One estimate was based on responses to our statutory information requests combined with IDC PaaS and IDC SaaS data. The other estimate was based on responses to our statutory information requests combined with Synergy PaaS and IDC SaaS data. We then divided total third-party sales through AWS, Microsoft and Google marketplaces by these estimates.
[296] Based on the revenue from first-party sales and commission fees from third-party sales. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

continue to grow.[297] Evidence suggests that the hyperscalers see marketplaces as an important distribution channel and are seeking to grow their use. Oracle noted that while marketplaces are not yet a major revenue source for hyperscalers, they "can act as particularly powerful generators of network effects", specifically for those customers who use marketplace spend to meet spend commitments.[298] [✂] aims for its marketplace to be the primary online distribution channel for both first-party and third-party solutions.[299] [✂] aims for its marketplace to become the most strategic channel that ISVs use to acquire new customers,[300] whereby [✂].[301]

*Co-selling and marketplaces are levers to drive underlying infrastructure consumption*

4.90 Evidence gathered from the hyperscalers suggests that they ([✂]) see co-selling schemes and marketplaces as ways to develop the ecosystem of services they offer on their infrastructure and ultimately as a lever to drive underlying infrastructure consumption. This is because ISV solutions bought through a particular provider's marketplaces will run on that provider's infrastructure. [✂] explained that ISVs are "an important sell-through channel" which are "actively selling solutions that run on [✂], helping us acquire new customers and driving more [✂] consumption from existing customers."[302] [✂] explained the critical importance of being a first-choice platform for ISVs as IT spend is shifting to software rather than infrastructure consumption.[303] [✂] explained that "marketplace is built primarily to drive underlying [✂] consumption",[304] and "marketplace partners will be offered increasingly tiered incentives and programs [✂] depending on how much they drive infrastructure use ([✂]) and also [✂] ([✂], [✂])."[305]

# Market outcomes

4.91 In this subsection we report key market outcomes, including the positioning of the major cloud providers in the market and we consider some positive outcomes highlighted in response to our interim report. We also look at the shares of supply for UK cloud services, analyse the profitability of the key providers, consider the service offerings and key capabilities of each hyperscaler, and the prevalence of multi-clouding in the market.

## UK shares of supply

### UK revenues for IaaS and PaaS

4.92 Table 4.8 summarises our estimates of UK revenues associated with IaaS and PaaS. We explain how we derived these estimates in Annex 1. We have focused on UK shares of supply rather than global shares, as global shares would likely understate the position of the

[297] For example, see Canalys Insights. Canalys Insights - Are cloud marketplaces worth the hype? [accessed 26 September 2023]. Also see [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[298] Oracle response to the MIR consultation, page 4.
[299] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[300] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[301] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[302] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[303] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[304] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[305] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].

hyperscalers in the UK, given the large Asian cloud providers, such as Alibaba, have a more limited presence in the UK.

4.93 We estimate that in 2022, cloud infrastructure services generated revenues of £7.0bn to £7.5bn. Between 2019 and 2022, UK revenues for IaaS and PaaS combined grew by 35% - 40% per year.

**Table 4.8: UK IaaS and PaaS revenues, £bn**

| | 2019 | 2020 | 2021 | 2022 | Annual growth |
|---|---|---|---|---|---|
| **IaaS** | [✂] [1.5-2.0] | [✂] [2.0-2.5] | [✂] [2.5-3.0] | [✂] [4.0-4.5] | 30-35% |
| **PaaS** | [✂] [0.5-1.0] | [✂] [1.0-1.5] | [✂] [1.5-2.0] | [✂] [2.5-3.0] | 40-45% |
| **IaaS and PaaS** | [✂] [2.5-3.0] | [✂] [3.5-4.0] | [✂] [4.5-5.0] | [✂] [7.0-7.5] | 35-40% |

*Source: Ofcom analysis of data provided in response to our information requests and data from Synergy and IDC. Annual growth based on the compound annual growth rate between 2019 and 2022.*

4.94 Our estimates indicate that UK IaaS revenues are higher than PaaS revenues, though PaaS revenues have generally been growing slightly quicker. However, there is more uncertainty around UK revenues for PaaS due to different estimates from the data sources available to us, as explained in Annex 1.

4.95 Based on data obtained from IDC, there are more providers of PaaS in the UK than of IaaS. While IDC does not track all providers, there could be in the region of 30+ IaaS providers in the UK and 200+ PaaS providers.[306]

4.96 The main types of service provided as IaaS are compute, storage and networking. Of these, compute typically generates the most revenue, and represented about two-thirds of UK IaaS revenue in 2022.[307]

4.97 Various types of service can be provided as PaaS. Based on our analysis of responses to our information requests, services associated with data management and analytics appear to represent most UK PaaS revenues for the hyperscalers in aggregate.

4.98 IaaS and PaaS revenues are expected to continue growing. Responses to our information requests referenced reports from industry analysts that forecast IaaS and PaaS revenues in the UK and Ireland could grow 25% to 30% per year until 2024/2025,[308] and this is broadly consistent with internal revenue forecasts received from cloud providers.[309] However,

[306] IDC, Public Cloud Services Tracker, 2022 H2 (published April 2023).
[307] Annex 1, Shares of supply for cloud infrastructure services in the UK.
[308] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[309] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

recent quarterly earnings releases from the largest providers indicate global revenue growth is likely to be slower in 2023.[310]

## UK shares of supply for IaaS

4.99 The table below shows our estimated shares of supply for UK IaaS. We explain how we derived these estimates in Annex 1.

4.100 Based on responses to our information requests, and the market size estimates set out above, we estimate AWS and Microsoft accounted for approximately [✂]% [80% to 90%] of UK IaaS revenues in 2022; a proportion that has slowly increased since 2019. Within this, Microsoft's share of UK IaaS revenues has grown while AWS's share has reduced slightly.

**Table 4.9: UK IaaS shares of supply, 2019 – 2022**

| | 2019 | 2020 | 2021 | 2022 |
|---|---|---|---|---|
| **AWS** | [✂]% [40-50%] | [✂]% [40-50%] | [✂]% [40-50%] | [✂]% [40-50%] |
| **Microsoft** | [✂]% [30-40%] | [✂]% [30-40%] | [✂]% [30-40%] | [✂]% [40-50%] |
| **AWS + Microsoft** | [✂]% [70-80%] | [✂]% [70-80%] | [✂]% [70-80%] | [✂]% [80-90%] |
| **Google** | [✂]% [0-5%] | [✂]% [0-5%] | [✂]% [0-5%] | [✂]% [0-5%] |
| **Other** | [✂]% [10-20%] | [✂]% [10-20%] | [✂]% [10-20%] | [✂]% [10-20%] |

*Source: Ofcom analysis of data provided in response to our information requests and data from Synergy and IDC. Some numbers may not sum due to rounding.*

4.101 While Google's UK IaaS revenues have grown since 2019, it represented [✂]% [0% to 5%] of UK IaaS revenues in 2022, significantly behind AWS and Microsoft.

4.102 The 'other' category includes a number of other providers with low IaaS shares of supply in the UK, and overall has steadily declined as Microsoft and Google have gained share. This category includes IBM and Oracle, with shares [✂]% [0% to 5%] in 2022. While UK IaaS revenues for smaller cloud providers grew between 2019 and 2022, [✂].

## UK shares of supply for PaaS

4.103 PaaS includes many diverse types of services. Many companies specialise in providing one type of service (e.g. data management services), while only a handful, like the hyperscalers, offer services across all PaaS categories.[311] The table below shows our estimated shares of supply for UK PaaS. We explain how we derived these estimates in Annex 1.

---

[310] In Annex 2 we analyse quarterly year-on-year global revenue growth for major cloud providers and show that cloud revenue growth has been declining over recent quarters for each of the hyperscalers, reflecting challenging macroeconomic conditions. For example, in its Q2 2023 earnings call, Amazon said that due to economic uncertainty over the last year, AWS customers "have needed assistance cost optimizing to withstand this challenging time and reallocate spend to newer initiatives that better drive growth". However, it also said that it has "started seeing more customers shift their focus toward driving innovation and bringing new workloads to the cloud". Amazon Q2 2023 earnings call (7 minutes 50 seconds to 8 minutes 30 seconds) [accessed 26 September 2023].

[311] By PaaS product category here we are referring to IDC's seven 'secondary markets' for PaaS which are analytics and business intelligence, AI platforms, data management, integration and orchestration, application development, software quality and life cycle, and application platforms. IDC, Public Cloud Services Tracker 2022 H2 (published April 2023).

**Table 4.10: UK PaaS shares of supply, 2019 – 2022**

| | 2019 | 2020 | 2021 | 2022 |
|---|---|---|---|---|
| **AWS** | [✂]% [20-30%] | [✂]% [20-30%] | [✂]% [20-30%] | [✂]% [20-30%] |
| **Microsoft** | [✂]% [10-20%] | [✂]% [20-30%] | [✂]% [20-30%] | [✂]% [20-30%] |
| **AWS + Microsoft** | [✂]% [40-50%] | [✂]% [40-50%] | [✂]% [40-50%] | [✂]% [50-60%] |
| **Google** | [✂]% [5-10%] | [✂]% [10-20%] | [✂]% [10-20%] | [✂]% [10-20%] |
| **Other** | [✂]% [40-50%] | [✂]% [40-50%] | [✂]% [30-40%] | [✂]% [30-40%] |

*Source: Ofcom analysis of data provided in response to our information requests and data from Synergy and IDC. Some numbers may not sum due to rounding.*

4.104 The table indicates that AWS and Microsoft represented [✂]% [50% to 60%] of UK PaaS revenues in 2022 – a lower share than for IaaS. Between 2019 and 2022, we estimate that Microsoft's share of UK PaaS revenues has grown slightly faster than AWS's share.

4.105 Google's share is closer to that of AWS and Microsoft in PaaS than in IaaS, with a [✂]% [10% to 20%] share of UK PaaS revenues in 2022. [✂].

4.106 The share of supply of the "other" category fell in each year between 2019 and 2022, and by around [✂]% [10% to 20%] points over the period. Oracle, MongoDB and IBM have some of the larger shares of supply of companies in the 'other' category; we estimate they represented around [✂]% [0% to 5%], [✂]% [0% to 5%] and [✂]% [0% to 5%] respectively of UK PaaS revenues in 2022.

### UK shares of supply for IaaS and PaaS combined

4.107 The table below shows our estimated shares of supply for UK IaaS and PaaS combined, drawing on the information presented above.

**Table 4.11: UK shares of supply for IaaS and PaaS combined, 2019 – 2022**

| | 2019 | 2020 | 2021 | 2022 |
|---|---|---|---|---|
| **AWS** | [✂]% [30-40%] | [✂]% [30-40%] | [✂]% [30-40%] | [✂]% [30-40%] |
| **Microsoft** | [✂]% [20-30%] | [✂]% [30-40%] | [✂]% [30-40%] | [✂]% [30-40%] |
| **AWS + Microsoft** | [✂]% [60-70%] | [✂]% [60-70%] | [✂]% [60-70%] | [✂]% [70-80%] |
| **Google** | [✂]% [0-5%] | [✂]% [5-10%] | [✂]% [5-10%] | [✂]% [5-10%] |
| **Other** | [✂]% [20-30%] | [✂]% [20-30%] | [✂]% [20-30%] | [✂]% [20-30%] |

*Source: Ofcom analysis of data provided in response to our information requests and data from Synergy and IDC. Some numbers may not sum due to rounding.*

4.108 We estimate that in 2022 AWS and Microsoft had around [✂]% [70% to 80%] share of UK combined IaaS and PaaS revenues, with Google significantly lower on [✂]% [5% to 10%]. Overall, we estimate that AWS, Microsoft and Google accounted for [✂]% [70% to 80%] of UK IaaS and PaaS revenues in 2022.

4.109 Between 2019 and 2022, the hyperscalers' overall UK share of supply of IaaS and PaaS combined increased. Over this period, we estimate that AWS's UK share of supply fell marginally, Microsoft's grew and Google experienced the strongest growth, although from a lower revenue base.

4.110 The 'other' category includes many companies with low shares, but IBM and Oracle appear to have some of the larger UK shares of supply in this category. We estimate they both represented around [✂]% [0% to 5%] of UK IaaS and PaaS combined revenues in 2022.

### UK shares of supply for SaaS

4.111 Compared to IaaS and PaaS, UK SaaS revenues are significantly more fragmented. There is much more diversity across SaaS services, market features and suppliers, and the segment is not characterised by the same level of concentration that we see in IaaS and PaaS.

4.112 In 2022, the hyperscalers' share of UK 'SaaS – Applications' revenue was around 18% according to IDC, most of which related to Microsoft services.[312]

## Positioning of the hyperscalers

4.113 In this subsection, we describe in more detail the positioning of AWS, Microsoft and Google, as the leading players in the provision of cloud infrastructure services in the UK. We set out the product categories and services that are particularly important for each of the hyperscalers, and the key capabilities they have developed which enable them to compete and attract customers. We refer to some of these capabilities in later sections, where relevant.

### AWS

4.114 AWS is generally recognised as offering a strong portfolio of services across both IaaS and PaaS. In 2022, AWS's public cloud infrastructure revenues were primarily driven by its [✂] and [✂] product categories.[313] Relative to the other hyperscalers, however, we note that AWS has fewer enterprise SaaS offerings, and instead its strategy appears more focused on enabling partners (e.g. ISVs) to build SaaS on top of AWS's infrastructure.[314]

4.115 AWS was the first to enter the cloud services market. AWS began by offering IaaS to customers, using the infrastructure it had built for its own retail business. As the first mover, AWS was therefore the only advanced provider present when some customers (e.g. [✂]) started to think about migrating to the cloud.[315]

4.116 AWS today is recognised to have the greatest breadth and depth of capabilities,[316] due to the range of services it offers, and the functionality offered within those services.[317] As noted above, AWS also [✂]. One source ([✂]) suggests that AWS positions itself as the best provider of cost effective, scalable infrastructure and platform as a service,[318] with another source suggesting that AWS may be particularly cheap for compute infrastructure due to its

[312] IDC, 2022. Public Cloud Services Tracker, 2022 H2 (published April 2023).
[313] AWS response dated 31 March 2023 to s.174 request of 24 October 2022, Part B, question 4.
[314] Computer Weekly, 2019. Should AWS consider building out its SaaS play? [accessed 9 February 2023].
[315] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[316] [✂].
[317] This includes its own cloud services, as well as third-party services (e.g. from ISVs).
[318] [✂] response dated [✂] to our proposed use of information dated [✂].

Graviton processor.[319] We also observe that BT Group, [✂] and [✂] consider AWS as a market leader in broadcast and video-processing, due to AWS's rich range of services to support this sector (e.g. Elemental technologies).[320]

## Microsoft

4.117 Microsoft is generally recognised to have strong capabilities across both IaaS and PaaS. In 2022, Azure's revenue was mainly driven by its compute and storage product categories.[321] However, customers also recognise Microsoft and/or may choose Microsoft for its enterprise SaaS offerings.[322]

4.118 Microsoft was well known for its provision of operating systems and productivity software before entering the cloud services market. Its position in traditional IT and SaaS makes Azure today particularly attractive for mid-size and large enterprises that are already using Microsoft's products.[323] Some customers that we engaged with identified Azure as a natural choice for such reasons. Our market research also suggests that among other reasons, 29% of Azure users chose Azure as their cloud provider due to already having an existing relationship with Microsoft for other services.[324]

4.119 Linked to this, one of Azure's key capabilities is the integration it can offer cloud customers with Microsoft's existing products. This was identified amongst respondents in our market research, with one respondent citing that "the integration with the other Microsoft systems is natural" when choosing Azure.[325] This ease of integration is also a feature of some of Azure's products. For example, Azure Active Directory is an enterprise identity service, which can enable employees within an organisation to sign into multiple services (across Microsoft 365 and Azure) with a single sign-on.[326] This may be particularly important for customers that require highly integrated features across their on-premises and cloud infrastructures (e.g. banks).

## Google

4.120 In 2022, Google Cloud's revenues were primarily driven by its [✂] and [✂] product categories.[327] Google Cloud is recognised to have particularly strong capabilities in PaaS, as reflected by its market share in PaaS being significantly higher than in IaaS. Google Cloud also offer a range of enterprise SaaS products (e.g. Gmail) within Google Workplace.[328]

4.121 Google is widely perceived to be a market leader in the provision of AI/ML and data analytics in cloud, with its provision of products such as BigQuery.[329] This specialism is recognised by a

---

[319] [✂]. The Graviton processor is used to power Amazon EC2 instance types. What is AWS Graviton? [accessed 26 September 2023].
[320] BT Group response to the CFI, page 19; [✂] response to the CFI, page [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[321] Microsoft response dated 25 April 2023 to s.174 request of 21 October 2022, Part B, questions 1 & 4.
[322] BT Group response to the CFI, page 19; and [✂] response dated [✂] to our customer questionnaire, question [✂].
[323] Datamation, 2023. AWS vs. Azure vs. Google Cloud [accessed 26 September 2023].
[324] Context Consulting research data tables, Q25.
[325] Context Consulting research report, slide 57.
[326] Microsoft website. Azure Active Directory [accessed 26 September 2023].
[327] Google response dated 31 March 2023 to the s.174 notice dated 26 October 2022, Part B, questions 1 & 4, Annex 2.
[328] Google website. Google Workspace [accessed 26 September 2023].
[329] BigQuery is a fully managed enterprise data warehouse with built-in features like machine learning, etc.

range of cloud providers and customers alike. Analytics is the only sub-category where [✂].[330]

4.122 Google Cloud have over time focused on designing services that encourage multi-cloud and hybrid-cloud environments. Google Cloud is recognised as one of the biggest adopters and promoters of open-source technologies,[331] having been responsible for first introducing Kubernetes, which enables the management of containerised applications. We observe that [✂].[332]

4.123 Google Cloud is also identified by other cloud providers for its aggressive pricing and intuitive user interfaces. Google suggest that in the first instance, they have to compete hard against competitors such as Oracle and IBM, to become a customer's secondary cloud provider, alongside either AWS or Azure (the primary cloud provider).[333] Cloud providers note that Google Cloud provide large discounts and generous credit offerings, to attract customers. Google Cloud is also recognised to have the ability to quickly set up customers, such as start-ups, on their infrastructure.[334]

## Key indicators of market outcomes – pricing and innovation

4.124 When considering broader market outcomes, some respondents to our interim report pointed to the existence of positive outcomes like continuous innovation and falling prices within the cloud market as evidence of strong competition.

### Pricing

4.125 Further to the findings in our interim report, where we noted that analyst reports suggested list prices of IaaS services have generally decreased over time,[335] AWS and Microsoft submitted analysis suggesting that list prices for these hyperscalers' core cloud services have fallen. While we have not undertaken our own analysis of pricing trends across the market, we have assessed the evidence in each submission and present our assessment below.

*AWS*

4.126 AWS provided data capturing how the UK list prices for three of its key services (S3, EC2 and data transfer-out (DTO)) have generally decreased over time.[336] AWS also provided global net price trend data for these services, which captures the overall average price that its customers have paid per unit of each service, taking account of discounts, refunds and actual customer use patterns.[337] There are some scope differences between the UK list price series

[330] Ofcom analysis of: [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[331] TechCrunch, 2019. Google remains the top open-source contributor to CNCF projects [accessed 27 September 2023].
[332] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response to [✂] of our follow-up email dated [✂] concerning the s.174 notice dated [✂].
[333] Google's response dated 23 November 2022 to the s.174 notice dated 26 October 2022, question 8.
[334] [✂] response dated [✂] to the s.174 notice dated [✂24 October 2022], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[335] Ofcom, 2023. Interim report, paragraph 6.12.
[336] [✂].
[337] [✂].

and global net price series (including geographical scope), which mean that these series are not perfectly comparable.[338]

4.127 We note that price trends for these three services are not necessarily representative of AWS's wider service portfolio, but these services are popular with AWS customers and [✂].[339]

4.128 AWS's data suggests that UK list prices for these three services have fallen over time. For S3 Standard and DTO from AWS Regions, UK list price reductions were concentrated between 2009 and 2016, with very few changes since 2017.[340] For EC2, the average UK list price across instance families fell by almost 30% between 2016 and 2022 (but a large price reduction between 2016 and 2017 means that prices fell by less than 15% between 2017 and 2022).[341]

4.129 AWS's global net price data indicates that the average global net prices of S3 and EC2 fell by [✂]% and [✂]% respectively between 2018 and 2022.[342] For DTO the average global net price fell by [✂]% between 2019 and 2022.[343]

4.130 Each of these reductions is larger than the corresponding UK list price reduction over the equivalent timeframe. We consider this tentatively suggests that price discounts have likely been increasing in importance, suggesting that larger customers (i.e. those that have the ability to negotiate and/or secure volume-based discounts) may have benefited more from price reductions than have customers who paid list prices. However, any comparisons between list price and net price trends are tentative due to the aforementioned differences in data scope. We also acknowledge that even where UK list prices have remained stable, the customers paying list prices might have benefited from other factors, such as improvements to the quality of services or the introduction of new service tiers (set at new list prices) that are better suited to their requirements than existing offerings.

4.131 AWS also provided data capturing how the global unit operating cost for each of these services has evolved relative to global net unit prices (over the same timeframes as the net price trends data). This data suggests that the global unit price for each of these services (calculated based on net prices) has fallen by a greater amount than global unit cost in recent years, resulting in a lower unit margin.[344]

---

[338] The differences in scope include: geography (global net prices versus UK list prices); product-level coverage (the global net price for S3 incorporates all S3 storage classes, whereas the S3 list price data relates to S3 Standard, [✂] (source: [✂]), and the global net price for DTO includes data transfers from AWS Regions to the public internet and data transfers from AWS Edge locations to the public internet via Amazon CloudFront, whereas the DTO list price data relates only to data transfers from AWS Regions to the public internet); time period (the time periods covered by each list price series and net price series vary, but for S3 and DTO the list price series date from 2008-09 to 2022 and are therefore significantly longer than the net price series, which date from 2017-18 to 2022-23).

[339] These three services collectively generated [✂]% of AWS revenue both in the UK and globally in 2022. EC2 was purchased by [✂]% of UK customer accounts in 2022, and S3 was purchased by [✂]% of UK customer accounts in 2022. Source: AWS response dated 31 March 2023 to the s.174 notice dated 24 October 2022, Part B questions 4 and 15.

[340] [✂].

[341] [✂].

[342] [✂].

[343] [✂]. For a general discussion on DTO, see: Joshua Gans, Mikaël Hervé & Muath Masri, 2023. Economic analysis of proposed regulations of cloud services in Europe, European Competition Journal, 19:3, 522-568, DOI: 10.1080/17441056.2023.2228668, pages 548-549.

[344] [✂]. [✂].

4.132 While this data suggests that unit margins are falling, we have not received detailed bottom-up evidence of how AWS has calculated these unit costs and unit margins. Moreover, we observe that unit costs [✂]. Therefore, despite AWS's evidence appearing to show that cost savings on these services have been passed onto prices, [✂].[345]

4.133 There is less evidence that AWS has [✂] as extensively, given the relatively gradual recent UK list price reductions noted above. This suggests that AWS has potentially [✂] some key services among customers paying UK list prices. However, we acknowledge that this finding is tentative, as we do not have unit cost data precisely tailored to the scope of the UK list price data AWS has provided.

*Microsoft*

4.134 Microsoft provided analysis capturing how the website (list) prices of its Azure IaaS and PaaS products have changed over time. Its data is global and incorporates website prices for a wide range of Azure products introduced since 2016, which collectively represent most of its IaaS and PaaS revenue. To account for the introduction of different products at different times, Microsoft's analysis measures the changes in product website prices for each month after a given product is introduced, and it aggregates these changes to form a revenue-weighted price index capturing average website price changes in the months after product introduction.[346]

4.135 Microsoft's data suggests that the revenue-weighted average global list (website) price of its Azure IaaS products decreases by around [✂]% [0%-5%] between the [✂] month after product introduction and the [✂] after product introduction. The equivalent decrease for Azure PaaS products is around [✂]% [0%-5%].[347] These are fairly marginal reductions in the price indices, and they do not account for the prevalence of price increases during the first [✂] months after product introduction.

4.136 Microsoft's analysis also captures the distribution of qualitative price changes across the product set it analyses. Specifically, this shows the proportions of total revenue associated with products which recorded no price changes, price increases, price decreases or both price increases and decreases over the 2016-23 period studied.[348]

4.137 This analysis shows that for a majority of Azure products by revenue (especially PaaS products), prices remained stable over the 2016-23 period. Without adjusting for increased quality of products, for IaaS products, price decreases were more common than price increases, but for PaaS products, price increases were slightly more common than price decreases. This appears broadly consistent with the findings of fairly marginal average price reductions in the price indices.

4.138 Microsoft's submission additionally provides a breakdown of global Azure Consumed Revenues (i.e. actual revenues) into implied list price revenue and the various discounts applied to these list prices (including committed spend discounts, reserved instances and saving plans, and credits) to generate Azure Consumed Revenue. This breakdown suggests that Azure Consumed Revenues have fallen as a share of implied list price revenue, from

[345] [✂]. This is consistent with our broader findings on AWS profitability, including AWS's persistently high returns, which we set out in the following subsection and in Annex 2.
[346] [✂].
[347] [✂].
[348] [✂]. [✂].

[✂]% in FY21Q1 to [✂]% in FY23Q3.[349] This implies that the global average net price charged for Azure services has fallen faster than the global average list price.

4.139 As with AWS, we consider this tentatively suggests that price discounts have likely been increasing in importance, suggesting that larger customers (i.e. those that have the ability to negotiate and/or secure volume-based discounts) may have benefited more from price reductions than have customers who pay list prices. However, we also acknowledge (as with AWS) that customers paying list prices might have benefited from other factors, such as improvements to the quality of services or the introduction of new service tiers. We also observe from Microsoft's submission that all customer groups (including PAYG) customers, who are not eligible for committed spend discounts) benefit from Reserved Instances and Saving Plans, which provide resource-specific discounts relative to list prices.[350] The importance of these discount types has increased in recent years.

*Our overall assessment of the evidence submitted by AWS and Microsoft*

4.140 Overall, the pricing evidence submitted by AWS and Microsoft suggests that list prices paid by PAYG customers for their core cloud infrastructure services have either remained stable or decreased in recent years. The data also suggests that average net prices for these services (which incorporate all types of discounts) have fallen faster than list prices over the same timeframe, although the list price and net price series are not always perfectly comparable. AWS has also told us that it has passed on global cost savings to customers through reduced prices, highlighting that the global unit margins for its core cloud infrastructure services have fallen. However, there is less evidence that AWS's cost reductions have been [✂].

## Innovation

4.141 [✂] told us that "cloud providers continue to innovate to differentiate their core services, and to pass on cost savings and efficiencies to customers through lower prices",[351] Google said that customers are "benefiting from product innovation, discounts and a wide choice of software services from ISVs".[352] Some customers also pointed to the consistent release of new innovations from cloud providers as a key benefit to using cloud or a driver in their choice of provider or services.[353] Microsoft pointed to evidence of innovation in its business and provided an illustrative list of new entrants that have contributed to innovation in cloud services in recent years.[354]

4.142 We believe that the scope for attracting new customers into a growing market is creating strong incentives for suppliers to innovate, and we recognise that continued innovation in the sector brings benefits to all customers. This goes for hyperscalers, as well smaller cloud providers and ISVs.

---

[349] [✂].
[350] [✂].
[351] [✂].
[352] Google response to the interim report, page 1, paragraph 3.
[353] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[354] Microsoft response to the interim report, pages 15-17; Microsoft response to the interim report (annexes), pages 7-8, paragraphs 27-33. CoreWeave (2017) – focused on Graphics Processing Unit computing, especially for generative AI technologies; Paperspace (2014) focused on generative AI workloads and serverless tool that abstracts from underlying hardware; DigitalOcean (2012) a self-serve SaaS business; Scaleway (2014) focused on providing cloud services to small businesses; Clever Cloud operates PaaS for developers and handles code deployment so developers can focus on coding.

4.143 However, as set out in more detail in Section 8, we are concerned that barriers to switching and multi-cloud could pose a risk to the extent of innovation in the future.

## Profitability

4.144 In this subsection, we:

a) compare hyperscaler operating profits and margins to those of other cloud providers; and
b) compare our estimates of return on capital employed (ROCE) for AWS and Microsoft's cloud businesses to the weighted average cost of capital (WACC).[355]

4.145 Our profitability analysis focuses on AWS, Azure and Google Cloud as the hyperscaler businesses providing cloud infrastructure services.[356] For Microsoft, we also reference Microsoft Cloud which, as well as Azure, includes Microsoft's other cloud services like Office 365 Commercial, the commercial portion of LinkedIn and Dynamics 365.[357] While Azure is a part of Microsoft Cloud, Microsoft Cloud's financial performance reflects the performance of all Microsoft's cloud activities, not just those related to cloud infrastructure.

4.146 We have run our profitability analysis at a global level because the major cloud providers are globalised businesses, with many of their expenses and investments in cloud services serving their global customer base.

4.147 We explain our approach to assessing profitability in more detail in Annex 2. In that annex we also respond to stakeholder comments on the profitability analysis we set out in our interim report. In general stakeholders did not challenge the detail of the EBIT and ROCE estimates we present below, and comments focussed primarily on the inferences drawn from these estimates.

### Hyperscaler operating profits compared to other global cloud providers

4.148 Amazon and Google publicly report earnings before interest and tax (EBIT) for AWS and Google Cloud respectively.

4.149 Microsoft does not publicly report EBIT for Azure or Microsoft Cloud.[358] Microsoft does publish revenues and gross margins for Microsoft Cloud, and it provided us with estimates of Azure's operating profit in response to our information requests. Using this data, we estimated EBIT for Azure and Microsoft Cloud. We explain how we did this in Annex 2.

4.150 Figure 4.12 shows the latest annual EBIT for AWS and Google Cloud, alongside our estimates for Microsoft Cloud and Azure (which has been redacted in the published version of this report) and compares this to the EBIT for other global cloud providers where EBIT data is available: Alibaba's reported 'Cloud' segment and DigitalOcean.[359] IBM and Oracle are not included in the chart as IBM provided only high-level information relating to cloud EBIT, and

[355] The definition of ROCE and WACC can be found in Annex 7.
[356] Google Cloud includes Google Cloud Platform, which provides cloud infrastructure services, and Workspace, which incorporates Google's consumer and enterprise SaaS like Gmail and Google Docs. While Google Cloud is broader than 'cloud infrastructure services', we think it gives a reasonable idea of Google's financial performance in cloud infrastructure for the purpose of this market study, as explained in Annex 2.
[357] Microsoft, 2023. 2023 10-K report, page 42 [accessed 26 September 2023].
[358] While Microsoft does publish EBIT for its Intelligent Cloud operating segment, this includes Azure alongside non-cloud server products, so it may not be representative of the profits associated with cloud infrastructure services.
[359] DigitalOcean provides cloud infrastructure services for startups and small and medium-sized businesses.

Oracle said it could not provide detailed or accurate estimates on cloud profit beyond its public reporting for the Cloud and License segment.[360]

4.151 Figure 4.12 shows that in absolute terms, the most recent annual EBIT for AWS and our estimate of EBIT for Microsoft Cloud are significantly higher than the EBIT of other cloud providers, including Google Cloud, which was consistently loss making until 2022 (although it reported its first quarterly operating profits in the quarters ending March and June 2023).[361,362] We estimate that Azure EBIT in Microsoft's 2023 financial year was £[✂]. This is [✂].

**Figure 4.12: Annual global EBIT for the latest financial year (£bn)**

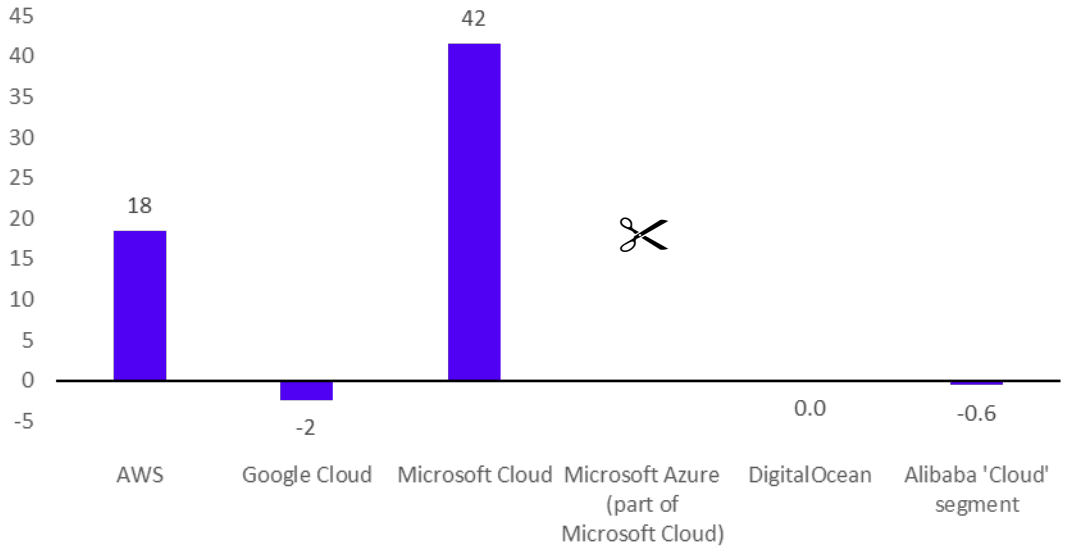


*Source: Ofcom analysis of cloud providers' published financial statements, information provided by Microsoft in response to our information requests[363] and Ofcom assumptions. Figures (other than Azure) come from the following financial statements: Microsoft – year to June 2023, AWS – year to December 2022, Google Cloud – year to December 2022, Alibaba 'Cloud' segment – year to March 2023, DigitalOcean – year to December 2022.*

4.152 Figure 4.13 shows quarterly EBIT margins for the same businesses between 2018 and June 2023. Our estimates for Azure have been redacted in the published version of this report.

[360] IBM response dated 23 December 2022 to our follow-up email dated 9 December 2022 concerning the s.174 notice dated 25 October 2022, Part B question 9; Oracle response dated 13 January 2023 to questions 7 and 14 of our follow-up email dated 22 December 2022 concerning the s.174 notice dated 31 October 2022, Part B questions 9 and 20; Oracle response dated 2 June 2023 to our follow-up email dated 10 May 2023 concerning the s.174 notice dated 31 October 2022, Part B question 9.

[361] Alphabet Inc., 2023. 2023 Q1 10-Q, page 29 [accessed 26 September 2023]; Alphabet Inc., 2023. 2023 Q2 10-Q, page 32 [accessed 26 September 2023].

[362] OVHcloud reports EBITDA for its 'Public Cloud' business. In its financial years ending 2021 and 2022 EBITDA was 35% to 40%. EBITDA is higher than EBIT as it is before depreciation and amortisation expenses. OVHcloud does not report EBIT margins for its Public Cloud segment, but EBIT margins for its overall business were close to zero in these years.

[363] Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).

**Figure 4.13: Quarterly EBIT margins for major cloud providers, 2018-23**

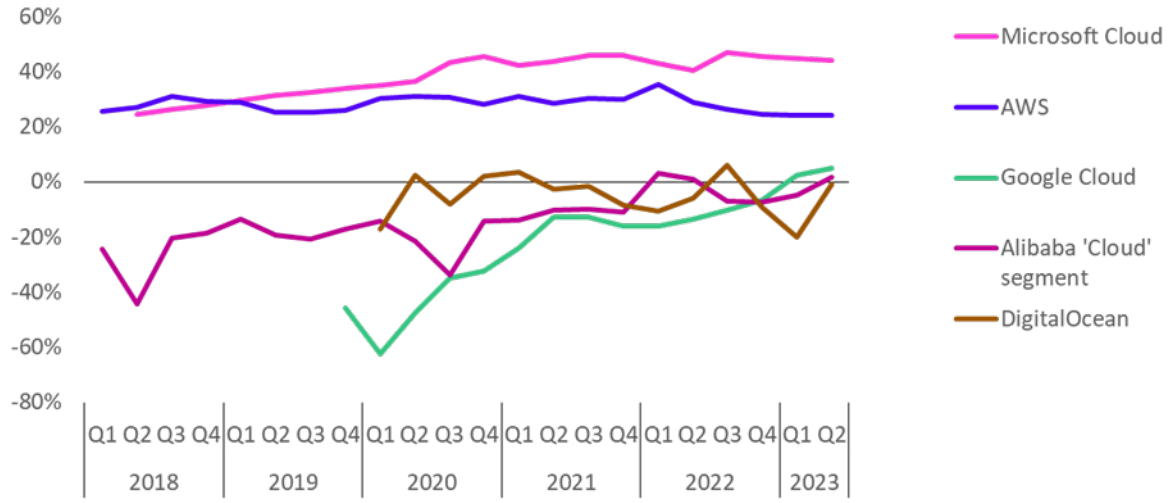


*Source: Ofcom analysis of cloud providers' financial data reported by S&P Capital IQ, information provided by Microsoft in response to our information requests[364] and Ofcom assumptions. We have used the latest filings (incorporating restatements).*

4.153 Figure 4.13 indicates that, over this period:

a) AWS had stable EBIT margins of around 20% to 30%.
b) Our estimated Microsoft Cloud EBIT margins increased from 25% to 45%. As Microsoft Cloud is broader than Azure, this data does not represent Azure's EBIT performance. Our estimated EBIT margin for Azure suggests [✂].
c) Google Cloud made losses until the end of 2022, but it has since reported operating profits for the first two quarters of 2023. Its EBIT margins are trending upwards but remain low compared to AWS and Microsoft Cloud.
d) Among other cloud providers, quarterly EBIT margins for Alibaba's 'Cloud' segment and DigitalOcean were occasionally positive but mostly negative in recent years.

4.154 Overall, this evidence indicates that cloud profits for AWS and Microsoft Cloud are higher than other cloud providers, for whom, in some cases, operating losses to date have been common. While there is some evidence of improving profits (or reduced losses) among smaller cloud providers, these are a lot lower than those for AWS and Microsoft Cloud. Azure operating profits [✂].

## Comparison of hyperscaler returns against WACC

4.155 Return on capital employed (ROCE) can be compared against the weighted average cost of capital (WACC) to assess how returns on investment compare to the cost of providing the capital to fund the business.[365] When combined with other indicators, a finding that ROCE is above WACC for a sustained period can be an indication of limitations in the competitive process. We say more about the use of ROCE to measure profitability in Annex 2.

4.156 Our ROCE analysis focuses on AWS and Microsoft Azure, as our shares of supply analysis indicates they represent a substantial share of cloud infrastructure revenues, and Google

[364] Microsoft response dated 20 January 2023 to our follow-up email dated 20 December 2022 concerning the s.174 notice dated 21 October 2022, Part B question 22; Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).
[365] ROCE is calculated by dividing EBIT by the value of capital employed in the relevant business.

Cloud has to date reported operating losses in its annual results (implying that its annual ROCE up to 2022 is negative). For context and comparison with Azure, we also present an estimate of ROCE for Microsoft Cloud.

4.157 We estimate that the pre-tax nominal WACC applicable to cloud services is likely to be between 9% and 13%. We explain this in more detail in Annex 2.

*AWS*

4.158 For AWS, our baseline ROCE is calculated based on the EBIT and net property and equipment assets for AWS reported by Amazon in its annual 10-K since 2013.[366] Our estimates are shown in the figure below.

**Figure 4.14: AWS ROCE estimates, 2013-22**

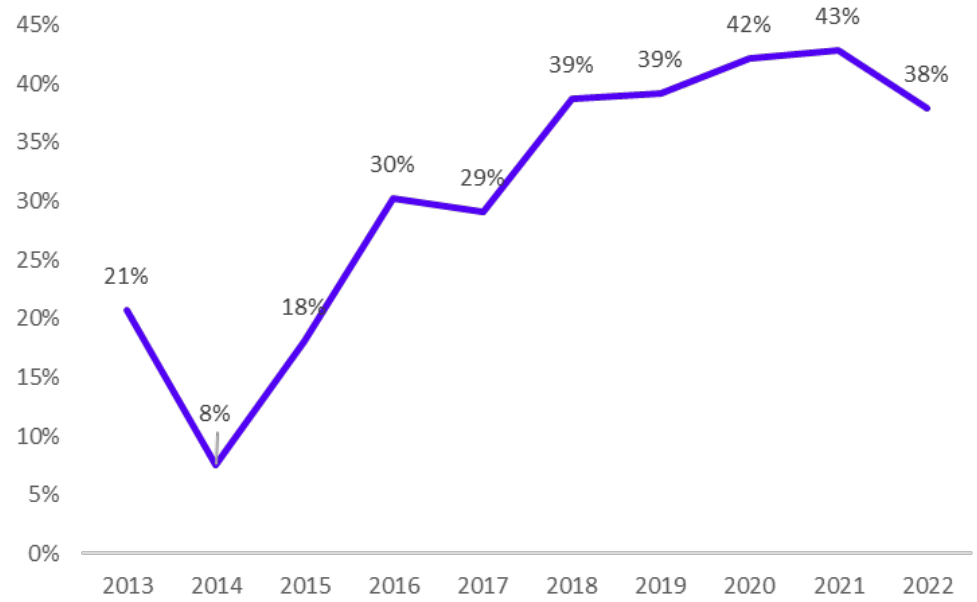


*Source: Ofcom analysis based on public information from Amazon 10-K reports.*

4.159 Our estimate indicates that AWS ROCE has increased since 2013 and has averaged 40% between 2018 and 2022. It is higher than our estimate of WACC in all years except 2014, i.e. in nine of the last ten years. Although not shown in Figure 4.14, we estimate that AWS ROCE reduced [✂]% in the year to June 2023. A fall in ROCE is consistent with recent quarterly reductions in EBIT, but it remains significantly above WACC.[367]

4.160 In Annex 2 we consider sensitivities to our AWS ROCE estimate, including taking account of additional assets and working capital within capital employed and attributing all Amazon's technology infrastructure assets to AWS. These sensitivities do not affect our analysis that AWS ROCE has been above WACC for a number of years.

*Microsoft*

[366] A 10-K form is an annual report required by the Securities and Exchange Commission in the US. It includes annual financial statements.

[367] Figure 4.13 shows that AWS's quarterly EBIT margins have fallen since early 2022. This follows a slowdown in revenue growth over this period as customers optimised their cloud spend in response to tough macroeconomic conditions. However, in its most recent Q2 2023 results announcement Amazon said "Our AWS growth stabilized as customers started shifting from cost optimization to new workload deployment".

4.161 Annex 2 details how we estimated EBIT and capital employed for Azure and Microsoft Cloud, for use in our ROCE calculation. Our estimates involve more assumptions than for AWS, as Microsoft reports less information on its cloud businesses, and it could not provide us with all the financial information we requested.

4.162 Our ROCE estimates for Azure for Microsoft's financial years ending 2018 to 2023 are shown in the figure below.

4.163 We estimate that Azure's ROCE increased over this period and was above our estimate of WACC for Microsoft's last three financial years. [✂].

**Figure 4.15: Azure ROCE estimates, Microsoft financial years ending 2018-23**

[✂]

*Source: Ofcom analysis based on Microsoft 10-K reports and information provided by Microsoft in response to our information requests.*[368]

4.164 In Annex 2, we consider whether the inclusion of intangible assets could impact our ROCE estimates for Azure, but conclude it is unlikely to change our findings regarding the trajectory and level of Azure ROCE.

4.165 Our baseline ROCE estimates for Microsoft Cloud for Microsoft's financial years ending 2016 to 2023 are shown in the figure below.[369]

**Figure 4.16: Baseline ROCE estimate for Microsoft Cloud, Microsoft financial years ending 2016-23**

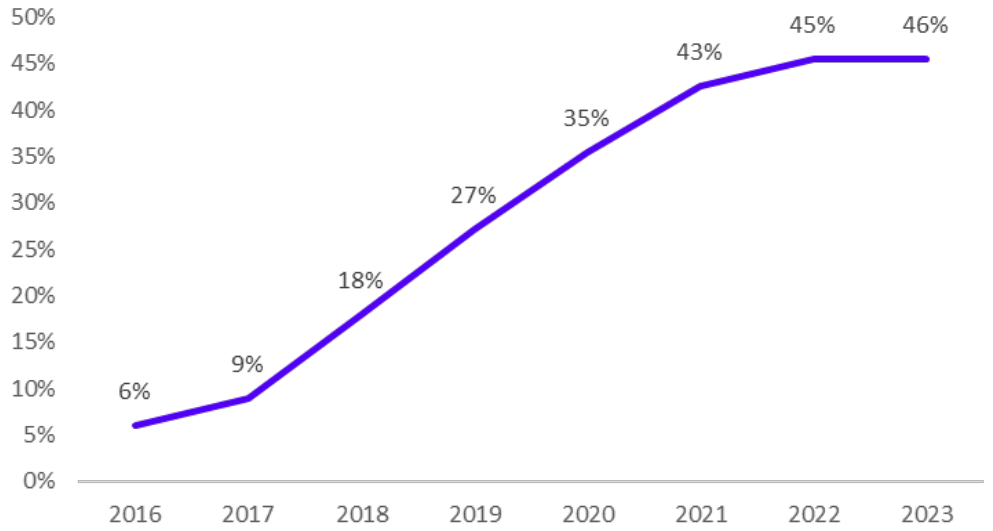


*Source: Ofcom analysis based on public information from Microsoft 10-K reports and Ofcom assumptions.*

4.166 We estimate that Microsoft Cloud ROCE steadily increased over this period. Our baseline estimate of Microsoft Cloud ROCE increased from 18% to 46% in Microsoft's financial years ending 2018 to 2023, above our estimate of WACC.

[368] Microsoft response dated 20 January 2023 to our follow-up email dated 20 December 2022 concerning the s.174 notice dated 21 October 2022, Part B question 22; Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).
[369] As noted in Annex 2, our estimate of ROCE for Microsoft Cloud was by reference to capital employed in Microsoft's overall business. This is likely to overestimate capital employed for Microsoft Cloud and underestimate ROCE.

4.167 In Annex 2 we consider a sensitivity to take account of working capital within capital employed. This reduces our estimate of Microsoft Cloud's ROCE, but not substantially enough to affect the observation that Microsoft Cloud ROCE appears to have been above WACC since at least Microsoft's 2019 financial year.

# 5. Barriers to multi-cloud and switching

5.1 In this section we consider the extent to which barriers to switching and multi-cloud may weaken effective competition and adversely affect market outcomes.

5.2 Customers may face four categories of barriers to switching and multi-cloud:

a) **Technical barriers.** These include: (i) technical efforts required to set-up and operate a preferred cloud architecture; and (ii) time and costs required to develop the skills needed to use different clouds.
b) **Egress fees.** These are financial charges customers face when they transfer data out from a cloud provider's infrastructure.
c) **Committed spend discounts.** These are discounts offered by cloud providers to customers who commit to purchase a minimum amount over an agreed period. Depending on how they are structured, they may present a commercial incentive for a customer to concentrate most or all of their spending with a single provider.
d) **Predicting cloud spend.** These are challenges associated with customers predicting their cloud spend, including price transparency and complexity of usage.

5.3 In the following subsections, we consider each barrier in turn, assessing the extent to which they may limit customers' ability to switch or multi-cloud and whether these obstacles may be particularly strong for specific use-cases or customer segments.

## Technical barriers

5.4 In this subsection we examine the extent to which technical barriers may hinder customers' ability to implement different multi-cloud architectures and switch.

a) First, we outline how technical barriers can hinder customers' ability to set up a multi-cloud architecture and switch.
b) Second, we assess whether a significant share of customers face high technical barriers to multi-cloud and switching.
c) Third, we set out the evidence, assessing whether some of these technical barriers may not be justified and might persist going forward.

5.5 Where relevant, we explain how technical barriers may also impact the ability for customers to build multi-vendor solutions. We then discuss the implications of this in Section 7.

### Technical barriers can hinder customers' ability to set-up a multi-cloud or multi-vendor architecture and switch

5.6 Customers rely on cloud infrastructure services to run applications and process data which they use to provide services internally and/or externally to their users. To ensure their cloud architectures work as effectively as possible, customers may wish to take-up cloud infrastructure services offered by different cloud providers (i.e. multi-cloud); mix first-party and ISVs' services within the same cloud (i.e. multi-vendor); switch some or all components

of their applications to a different cloud (i.e. switching between clouds); or, to an ISV hosted on the same cloud (i.e. switching within clouds).

5.7 As further detailed in Annex 4, the degree of interoperability and portability offered by cloud providers and ISVs are crucial for customers' ability to deploy their preferred cloud architecture and modify it when their preferences and technical requirements evolve. If a customer wishes to deploy a multi-cloud or multi-vendor architecture, a high degree of interoperability facilitates the integration of cloud infrastructure services offered by different cloud providers or ISVs.[370] If a customer wishes to switch between or within clouds, the degree of interoperability and portability will affect the scale of the technical efforts needed to successfully carry out the change.[371]

5.8 AWS and Microsoft said that the introduction of cloud services has materially improved interoperability and portability of services, such that barriers to switching are considerably lower compared to legacy IT services.[372] We recognise that IT markets have always been prone to some technical barriers and that cloud has in some ways supported greater interoperability and portability than before. However, in our market study, our aim is to consider the features of the market today and assess whether technical barriers have, or may have, adverse effects.[373]

5.9 We consider that customers may face several technical barriers which can reduce interoperability and portability in the cloud and, as a result, hinder their ability to implement different multi-cloud or multi-vendor architectures and switch. These technical barriers can be categorised as follows:

a) **Technical differentiation of cloud infrastructure services.** Cloud providers and ISVs offer different cloud infrastructure services. Technical differences at interface level (e.g. proprietary APIs, protocols and workflows) can reduce interoperability, while technical differences at functionality level can reduce portability. The greater the degree of such technical differentiation, the more effort is required of customers to deploy a multi-cloud or multi-vendor architecture and switch between and within clouds.
b) **Technical differentiation of ancillary cloud services.** Cloud providers and – to a more limited extent – ISVs offer different ancillary cloud services to support the operationalisation of their cloud infrastructure services, including observability, billing, security[374] and orchestration services. Technical differentiation between ancillary cloud services of different clouds or ISVs increases the complexity of managing a multi-cloud and, to a lesser extent, multi-vendor environment. This is because customers would need to use and continuously sync different sets of ancillary tools or sustain material technical efforts to create a unified interface. These barriers to multi-cloud or multi-

[370] As discussed below, the extent of a multi-cloud integration is on a spectrum, with some types of integrated multi-cloud requiring tight integration and therefore a high level of interoperability and portability.
[371] As discussed below, a lack of interoperability and portability require customers to materially change their application before it can run on the new architecture. Moreover, customers wishing to switch between clouds will typically port small parts of their architecture over time (e.g. single workloads or individual services) and so will need to set-up an integrated multi-cloud in the interim which would be more challenging where there is a low degree of interoperability.
[372] AWS response to the interim report, paragraph 9; [✂]; Microsoft response to the interim report, paragraph 179.
[373] Section 130A(2) of the Enterprise Act 2002.
[374] In particular, Identity Access Management (IAM). IAM services enable organisations to manage digital identities and control user access to critical corporate information within their could architectures.

vendor solutions can in turn hinder customers' ability to switch between and, to a lesser extent, within clouds.

c) **Integration and operationalisation efforts.** The technical efforts that go into the integration and operationalisation of cloud architectures may be allocated to different parties depending on the type of cloud architecture and the identity of the supplier of cloud services. In a single-cloud architecture, first-party cloud infrastructure services are pre-integrated by the cloud provider. In a multi-vendor architecture, integrating ISVs' services may require some technical efforts by the customers, the extent of which will typically depend on the amount of pre-integration that the ISV can achieve. In a multi-cloud architecture, the technical efforts typically sit with the customer. This makes it more difficult for customers to set up and operationalise their cloud architectures when they span multiple clouds and, to a lesser extent, when they combine first- and third-party cloud infrastructure services hosted on the same cloud. This can in turn hinder customers' ability to switch between and, to a lesser extent, within clouds.

d) **Asymmetry of functionalities.** First-party cloud infrastructure services may offer fewer functionalities when used in combination with services hosted on a different cloud or third-party cloud services hosted on the same cloud. This may discourage customers from setting up a multi-cloud or multi-vendor architecture and, as a result, also hinder their ability to switch between or within clouds.

e) **Data gravity.** The cloud where the bulk of a customer's data is hosted is likely to attract more of this customer's data, as well as associated customer applications and cloud services. This is because customers find it simpler to co-locate data, applications and cloud services to mitigate external challenges, such as latency, data governance and data sovereignty. This may further discourage customers from deploying a multi-cloud architecture or from switching between clouds. On the other hand, data gravity would not apply to multi-vendor architectures as in those cases first-party services and ISVs services would be hosted on the same cloud.

f) **Lack of technical skills.** The technical differentiation of clouds means a different set of technical skills is needed to work with different clouds. This requires customers to develop specific skills for each cloud they would like to use.[375] Lack of technical skills can therefore particularly add to the technical efforts required to multi-cloud and switch between clouds.

g) **Lack of transparency.** A lack of transparency about any of the above dimensions may further exacerbate the technical efforts required of customers to work with multiple clouds, multiple vendors and switch between and within clouds.

5.10 For ease of reference, the table below summarises the extent to which each of the above barriers might in principle hinder customers' ability to work with multiple clouds or vendors and switch between them.

[375] As detailed below, a multi-vendor cloud architecture combines first- and third-party services hosted on the same cloud. In such scenario, skills challenges are likely to be relatively low as customers would not need to develop a fundamentally new set of skills when setting-up or switching to a multi-vendor architecture.

**Figure 5.1: Summary of impact of technical barriers on customers' ability to integrate and switch**

| Technical barrier | Impact of technical barrier on customers wishing to: | | | |
|---|---|---|---|---|
| | Multi-cloud[376] | Switch between clouds | Multi-vendor | Switch within clouds |
| **Tech differentiation of cloud infrastructure services** | Strong | Strong | Strong | Strong |
| **Tech differentiation of ancillary cloud services** | Strong | Strong | Moderate | Moderate |
| **Integration and operationalisation efforts** | Strong | Strong | Moderate | Moderate |
| **Asymmetry of functionalities** | Strong | Strong | Strong | Strong |
| **Data gravity** | Strong | Strong | Low | Low |
| **Lack of technical skills** | Strong | Strong | Low | Low |
| **Lack of transparency** | Strong | Strong | Strong | Strong |

*Source: Ofcom.*

5.11 In practice, the overall scale of technical efforts required of customers will typically be a result of a combination of the above challenges in relation to a specific multi-cloud or switching scenario. Therefore, the following sections present evidence to assess the impact and nature of such barriers in the round and refer to specific barriers where relevant.

## The evidence suggests that a significant share of customers face high technical barriers to multi-cloud and switching

5.12 The evidence we have received throughout the course of this market study indicates that a significant share of customers is likely to face high technical barriers to multi-cloud and switching. In this subsection, we discuss this evidence as follows:

a) First, we set out evidence from the Context Consulting market research indicating that many customers cite technical challenges as barriers to switching and multi-cloud.
b) Second, we present evidence suggesting that developing cloud-specific skills can add to the technical barriers faced by customers.
c) Third, we explain why technical challenges are likely to be particularly acute for more integrated forms of multi-cloud and for switching between clouds.
d) Lastly, we present evidence indicating that hyperscalers' services that facilitate multi-cloud and switching are not sufficient to meaningfully mitigate technical barriers.

[376] As discussed below, the evidence suggests that other things being equal, these barriers are likely to materially hinder customers' ability to set-up more integrated forms of multi-cloud.

## Many customers cited technical challenges as a barrier to multi-cloud and switching

5.13 The Context Consulting market research found that 52% of customers cited a lack of interoperability between different IaaS/PaaS providers' services[377] as a key concern with the cloud infrastructure services market.[378] This level of concern is broadly consistent across users of all the major IaaS/PaaS providers, sectors and category of service used (i.e. IaaS and PaaS). The level of concern is higher for early adopters (58%), those who have added an IaaS/PaaS provider (56%) and those who have switched (60%). These groups of customers may provide a better indication of the importance of technical challenges as they are likely to have more practical experience with interoperability in the market.

5.14 With regards to multi-cloud, at least half of IaaS/PaaS users in the market research only use a single cloud provider.[379] This is likely due to technical barriers, at least in part. The market research found that around 65% of customers cited at least one of the following technical challenges as a barrier to using multi-cloud: interoperability challenges, technological challenges and lack of skills. We note that the top two challenges of using multi-cloud were moving data across IaaS/PaaS providers (45%) and greater costs/less cost efficiency (34%).[380] These are likely to capture some additional technical challenges. For example, moving data across IaaS/PaaS providers may be challenging because of low data portability, and greater costs may be related to increased technical effort in maintaining an additional abstraction layer in the cloud stack and subscription fees for adaptors.

5.15 The market research found that around 70% of customers cited at least one of the following technical challenges as a barrier to completely switching their IaaS/PaaS provider: interoperability challenges, application portability challenges, data portability challenges and the need to retrain staff. Moreover, 37% of customers cited at least one of those technical challenges as their main barrier to completely switching their IaaS/PaaS provider. We note that the most cited barrier to completely switching the IaaS/PaaS provider was time and cost of making the change (43%).[381] This is likely to capture some additional technical challenges, such as the time and costs needed to reconfigure applications due to a lack of interoperability. Indeed, the feedback we have received from customers suggests that in some cases they may consider the reconfiguration effort as time and cost consuming rather

[377] We note that the Context Consulting market research referred to IaaS/PaaS providers. Customers may have understood such providers to include both cloud providers and ISVs. Hence, some of the market research results may need to be interpreted with caution. Potential caveats to the market research results are noted where relevant in this section and set out in more detail in Annex 3.

[378] Context Consulting research data tables, Q63. The figure is 51% after excluding respondents using only private cloud.

[379] 52% of IaaS/PaaS users reported using more than one IaaS/PaaS provider. This is likely to overstate the use of more than one public cloud provider. This is because, in addition to using multiple public clouds, some respondents who use more than one IaaS/PaaS provider may be combining: (i) the products of an ISV and public cloud provider on the same cloud; (ii) private and public cloud solutions (i.e. hybrid cloud); or (iii) two private cloud providers. We discuss this further in Annex 3.

[380] Context Consulting research data tables, Q31. After excluding respondents using only private cloud, the relative frequency of the top two barriers is: moving data across IaaS/PaaS providers (45%) and greater costs/less cost efficiency (35%).

[381] Context Consulting research data tables, Q52 and Q53. This figure is 43% after excluding respondents using only private cloud. This was also the most cited option by all respondents, when asked to identify the single main barrier to completely switching the IaaS/PaaS provider (20%). This figure is 19% and maintains its relative ranking after excluding respondents using only private cloud.

than technically difficult (i.e. they know how to reconfigure their applications, but it would take time and money).

## Developing cloud-specific skills can add to the technical barriers faced by customers

5.16 Different cloud providers and ISVs use different proprietary cloud technologies (e.g. APIs, software development kits, protocols, workflows, programming languages and data formats). This requires customers to develop specific skills for each cloud provider or ISV they would like to use, which can add to the technical efforts required to switch and set up a multi-cloud or multi-vendor architecture. For example, cloud developers or engineers may have become proficient in using a particular provider's workflows, tools and APIs. If the customer wanted to switch to a new provider, it would need to invest time and resources into retraining or hiring new staff to be able to operate in that new environment.

5.17 The impact of such skills challenges is likely to be much lower in a multi-vendor environment since customers would not need to change the cloud environment they operate within. Setting up or switching to a multi-vendor architecture can usually be achieved without a change of personnel or material additional training. As such, in the following paragraphs we discuss skills challenges only in relation to multi-cloud and switching between clouds.

5.18 There is a significant degree of skills specialisation in the cloud. For example, the formal training courses and accompanying certification offered by each of the hyperscalers are specific to their individual cloud services. So cloud technical staff would need to obtain separate certifications to demonstrate expertise in an equivalent suite of services from multiple providers. One respondent to the market research said that the cloud environments of AWS and Microsoft differ to an extent such that the skills of technical staff are specific to one cloud.[382] A customer ([✂]) provided anecdotal evidence that, for the most part, workforce tends to specialise towards one cloud provider and may not be interested in retraining for another.[383]

5.19 Companies looking to switch to or add a new cloud provider will face costs associated with retraining existing staff or hiring new staff. The market research found that 'the need to retrain staff' was the second most cited challenge to switching provider, with 33% of respondents perceiving it as a barrier to switching and 8% considering it the most important barrier. This was raised as a concern more often by customers in the public sector (45%) or healthcare industry (43%). Moreover, 26% of participants in the market research perceived staff resistance to change as a potential challenge of switching provider. This may be due to preferences for working in a particular cloud environment, or perhaps because technical staff have a better understanding of the disruption that retraining or switching more generally would cause. Consistent with this, the responses to our customer questionnaire indicated that the need to retrain or hire new staff can add to the effort required to switch or add a cloud provider.[384]

5.20 Skills challenges can be exacerbated by periods of labour market shortages in the tech sector, with customers competing against global tech giants (including the hyperscalers) for talent. One customer ([✂]) told us that if it wanted to switch provider or adopt multi-cloud, the cost to train or hire staff would be especially expensive due to a world-wide shortage of

[382] Context Consulting research report, slide 118.
[383] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[384] See Annex 3.

skilled software engineers.[385] This may be more acute in sectors which require higher degrees of technical expertise, such as telecoms, due to the inherent complexity of the technologies involved. We recognise that the tech sector is currently going through a cycle of layoffs, which may alleviate the hiring challenges caused by skills shortages, though there is evidence to suggest the impact on cloud was minimal.[386] However, regardless of labour market conditions, customers may find it difficult to find technical staff specialised in the particular cloud they are considering switching to or adding as they are competing for skilled labour against large tech companies, who likely have more resources available to them to attract talent.

5.21 The hyperscalers provide some free training and are making efforts to improve their training offerings, [✂],[387] which could reduce the retraining costs of a company looking to switch providers. For example, Microsoft has sought to consolidate its training offerings into programmes available to all customers through its Enterprise Skills Initiative, which gives access to training to smaller customers who previously may not have been eligible.[388] [✂].[389] [✂].[390]

5.22 While the training courses offered by hyperscalers may lower skills challenges, they do not remove the opportunity costs in terms of time and effort to retrain their workforce. Some skills may be transferable across clouds and customers may train their staff to work across clouds irrespective of whether they intend to use a different or additional cloud. However, overall, based on our survey results and the qualitative feedback received by customers we consider that developing cloud-specific skills is still likely to add to the technical barriers faced by customers.

### The evidence suggests that integrated forms of multi-cloud and switching between clouds are more challenging to implement

5.23 In the interim report we examined the technical challenges customers may face in relation to different multi-cloud architectures (integrated, duplicated and siloed multi-cloud) and different switching scenarios (switching between and within clouds). We continue to consider that this categorisation of multi-cloud and switching provides a useful benchmark to examine technical challenges customers may face.

5.24 The evidence we have received since the interim report has clarified that there is not a clear-cut separation between different multi-cloud architectures. Instead, there is a spectrum of possible implementations going from the ones with little to no integration (i.e. the more 'silo-ed' end of the spectrum) to the ones with various integrations between cloud services, customers' applications and data (i.e. the more 'integrated' end of the spectrum).[391]

---

[385] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; also see Forbes, 2021. Is There A Developer Shortage? Yes, But The Problem Is More Complicated Than It Looks [accessed 20 September 2023].
[386] See, for example CRN, April 2023. AWS Confirms Layoffs Impacting 'Single Digit Percentage' Of Employees [accessed 13 September 2023].
[387] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[388] Microsoft blog, 2020, How Microsoft helps customers adopt Azure through developer education [accessed 28 September 2023].
[389] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[390] [✂].
[391] An extreme version of integrated multi-cloud is where applications dynamically distribute microservices between different clouds depending on a predetermined set of rules (e.g. which cloud offers the cheapest storage).

5.25 The evidence has also clarified that integrated multi-cloud is a necessary step in fully switching between clouds. This is because customers wishing to fully switch between clouds will typically migrate small parts of their architecture over time (e.g. single workloads, individual services) and will need to set-up a temporary integrated multi-cloud spanning the origin and target clouds.[392]

5.26 In light of this, the following paragraphs discuss: (i) technical barriers customers may face when implementing more integrated forms of multi-cloud; and (ii) technical barriers customers may face when switching between clouds. A more comprehensive overview of the technical challenges relating to individual multi-cloud architectures and switching scenarios is set out in Annex 4.

5.27 Overall, in line with our preliminary conclusion in the interim report, using multiple clouds is always likely to require some technical effort, but its scale is likely to be a function of the desired level of integration between clouds. Similarly, switching is likely to require material effort, but its scale will depend on the actual switching scenario and use-case.

*Technical barriers to integrated multi-cloud*

5.28 Integrated multi-cloud is where customers build their public cloud architecture by mixing and matching cloud services hosted on different public clouds. For an integrated multi-cloud deployment to work effectively, different customer applications, customer data and cloud services hosted on different clouds need to closely interoperate with each other. Hence, technical challenges mainly revolve around limitations to cross-cloud interoperability.

5.29 Our evidence highlights technical challenges (as well as financial costs) associated with this type of deployment. In particular:

a) The qualitative part of the market research found that for most customers integrated multi-cloud is the desired model, but the challenge of making multiple clouds work in an integrated way is an obstacle, especially for larger organisations. Lack of interoperability was most commonly cited as a significant obstacle and usually stems from the difficulties of making one cloud stack work with another (particularly in the case of Azure). A minority of respondents said they have not experienced significant obstacles to a partially integrated multi-cloud set-up, but these companies tend to be smaller and have simpler technical requirements.[393]
b) The responses to our customer questionnaire indicate that, due to low interoperability, integrating multiple public clouds would require considerable work which has a significant cost and does not allow customers to focus on the areas where their business adds value.[394]

---

[392] Depending on how complex the customer's architecture is, the integration of the origin and target clouds may last for weeks, months or even years.
[393] Context Consulting research report, slide 80.
[394] [✂] cited limitations to or lack of interoperability as one barrier to integrated multi-cloud ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); [✂] ranked the need to materially reconfigure data and applications as the second most important barrier to integrated multi-cloud after nature of workloads or applications to be run across cloud providers ([✂] response dated [✂] to our customer questionnaire, question [✂]); [✂] said that the limitation of interoperability is the main barrier to integrated multi-cloud ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); [✂] said there is lack of interoperability and/or difficulty in incorporating services from multiple providers in a single cloud architecture especially at the speed ([✂]); [✂] cited limitations to or a lack of interoperability as one of the main barriers to integrated multi-cloud ([✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]); [✂] said

c) Some smaller cloud providers mentioned that a lack of interoperability between clouds can act as a barrier to integrated multi-cloud and hinder their ability to acquire customers.[395]

d) The existence of technical challenges to integrate multiple clouds has generally been confirmed in the evidence we have received since the interim report. For example, a customer ([✂]) said that it doesn't necessarily want to split a single function across clouds due to the complexities of managing two incompatible IaaS stacks. In particular, the customer said that integration within a hyperscaler ecosystem is an issue for cross-cloud interoperability or for setting-up distributed data or applications.[396] [✂] explained that in many cases integrating 'best of breed' capabilities from multiple clouds is technically feasible but there are additional complexities to this way of working which entail additional work and costs on the customer side.[397]

e) As further detailed in Annex 4, our analysis of hyperscalers data indicates that, across all hyperscalers, customers spending more than £10k per year (who account for a significant portion of the hyperscaler's revenues) use on average at least [✂] [10 to 20] proprietary cloud services within that hyperscaler's cloud. Our analysis suggests that customers using this number of proprietary services would likely face material technical barriers if they wanted to replicate some or all parts of such an architecture across multiple clouds.

5.30 We recognise that some integration between clouds may occur in certain cases. However, the evidence we have seen indicates that, where it occurs, multi-cloud is predominantly on the siloed end of the spectrum and this is partly because it is technically less demanding:

a) According to the market research, siloed multi-cloud is the most frequently adopted architecture by customers using multiple IaaS/PaaS providers (45%).[398] However, this

---

that integrated multi-cloud requires the use of an additional abstractions layer which risks access to innovation, and leads to higher development costs and worse performance ([✂] response dated [✂] to our customer questionnaire, question [✂]); [✂] said that the main barrier to integrated multi-cloud is the effort to configure two different environments ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); [✂] ranked limitations to or lack of interoperability as the top barrier to integrated multi-cloud ([✂] response dated [✂] to our customer questionnaire, question [✂]); [✂] said complexity in terms of interoperability has been a key factor in it not having adopted integrated multi-cloud ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]).

[395] BT Group said that the practicality of adopting a multi-cloud strategy might be limited because of interoperability concerns (BT Group response to the CFI, Appendix A10); BT Group also said that cloud services differ in all aspects from one provider to another, even if they offer broadly equivalent services, which limits interoperability in most cases and acts as a barrier to integrated multi-cloud (BT Group response dated 27 January 2023 to the s.174 notice dated 1 December 2022, question 18); Cloudflare said that the improvement to interoperability would help foster a multi-cloud system (Cloudflare response to the CFI, page 2); Oracle said that true multi-cloud means that a customer can mix, match, interconnect, and interoperate among all the varied cloud providers' services, and it said that multi-cloud is procompetitive and benefits all customers but UK customers cannot successfully achieve a multi-cloud strategy where switching costs are kept artificially high by market participants with significant market power (Oracle response to the MIR consultation, page 3); [✂] said that it believes that the cloud industry currently faces several technical barriers that limit competition on the merits, and that these barriers arise from restraints to interoperability and portability, which are critical to avoid vendor lock-in and reduce switching costs between the different cloud platforms ([✂] response dated [✂] to the s. 174 notice dated [✂], pages [✂] and questions [✂]).

[396] Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂].

[397] Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂].

[398] This figure is 47% after excluding respondents using only private cloud (Context Consulting research data tables, Q29). As discussed in Annex 3, 40% of customers (39% excluding users of private cloud only) who 'spread similar workloads across clouds' also may capture some siloed multi-cloud.

may not be a result of strong customer preferences for siloed multi-cloud and may simply reflect the relative ease of adopting this approach compared to other models.[399] Indeed, the qualitative part of the market research found that customers would generally prefer to adopt an integrated multi-cloud architecture, but this is deemed difficult to implement.[400] This suggests that in some circumstances siloed multi-cloud may be chosen as an easier (but unpreferred) multi-cloud architecture, from customers wanting to procure specific services from different cloud providers.

b) This was generally confirmed in our engagement with customers. For example, [✂] described their architecture as 'loosely coupled' and said that it has made the decision to concentrate on AWS in order not to deploy any additional abstraction layers.[401] Two customers ([✂] and [✂]) explained that they have adopted a siloed multi-cloud architecture as running workloads across clouds would be inefficient[402] or impose additional technical or commercial costs.[403] Another two customers ([✂] and [✂]) reported using a siloed multi-cloud architecture and did not cite any particular technical challenges with adopting this multi-cloud approach.[404]

c) Evidence from suppliers of professional services indicates that integration between clouds is usually aimed at creating some connection between siloed applications and is limited in scope.[405] In particular, [✂] suggested that customers typically ask to create a multi-cloud strategy with some connection between siloed applications and operational resilience recovery abilities, rather than a cloud architecture leveraging services across clouds to deliver a single application.[406] [✂] said that it does not see much integration of multiple public clouds from a functional perspective and assumed that this may be because: (i) customers are in the early stages of their cloud adoption, which means their focus is on migration to the cloud and integration from an operational perspective whereas they lack the required skillset for integration from a functional perspective; (ii) data gravity means that it's easier to do anything where data sits because the data is there and that environment has already been operationalised.[407]

d) A [✂] report [✂] indicates that it is typical for multi-cloud organisations to concentrate 80% or more of their workloads with their primary strategic provider.[408]

5.31 We have also received evidence indicating that several stakeholders see a range of benefits from integrated multi-cloud, indicating that there is an interest and potential demand for more integrated forms of multi-cloud:

a) Some customers highlighted multiple benefits from using an integrated multi-cloud. For example, [✂] told us that it has appetite for an integrated multi-cloud partly to gain

---

[399] The market research indicates that 34% of those who run different applications on different clouds see a lack of interoperability between clouds as a challenge to using multiple cloud providers. Context Consulting research data tables, Q31.
[400] Context Consulting research report, slide 80.
[401] Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂].
[402] [✂] response dated [✂] to our customer questionnaire, question [✂].
[403] [✂] response dated [✂] to our customer questionnaire, question [✂].
[404] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[405] We note that suppliers of professional services often act as system integrators, whose role is to help customers integrate multiple clouds. The view of these suppliers on the current prevalence of more integrated forms of multi-cloud may not be representative of the full market.
[406] Ofcom / [✂] meeting, [✂].
[407] Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂].
[408] [✂]

bargaining power from its ability to more easily switch between clouds.[409] [✂] said that, if mixing and matching best-of-breed capabilities or services from different clouds was easier, that may be advantageous, but it is difficult to predict how much it would make use of different capabilities as it is still relatively new into its cloud journey.[410] [✂] said that it believes in using multi-cloud for an end-to-end IT service (i.e. to integrate functions hosted on different clouds), if there are benefits for that customer.[411] Other customers ([✂] had a more neutral position – they did not anticipate any growth in demand for integrated multi-cloud but also suggested that there are barriers that prevent this.[412]

b) Some smaller cloud providers highlighted the importance of integrated multi-cloud growing in the future. Oracle said that the end-state of true multi-cloud means that a customer can mix, match, interconnect, and interoperate among all the varied cloud service providers' services. It said that, as the technology moves forward, every cloud provider will have a different approach on how to design its cloud architecture and it is important for customers to be able to access their innovation through easy integration of multiple-clouds and switching.[413] IBM said that it sees the cloud as a technology rather than as a physical place. IBM said that this is why it deploys many of its individual cloud services on other clouds (e.g. AWS and Microsoft). This allows customers using those other clouds to easily create a mixed architecture by combining first-party proprietary services with IBM services hosted on those clouds.[414]

c) In a survey conducted by Public First, respondents that used more than one cloud were asked about the current level of integration between their different cloud providers and how important they considered the following types of integration: application integration, management integration, security integration and data integration. While the majority of respondents said that integration between different cloud platforms was 'very important' or 'somewhat important' for all types of integration, only 10% said their use of different cloud providers is currently largely integrated.[415]

5.32 Overall, technical barriers to deploy an integrated multi-cloud architecture are likely to be material and a function of the desired level of integration between customers' applications, data and cloud services hosted on different clouds.

*Technical barriers to switching between clouds*

5.33 Switching between clouds involves customers migrating their cloud architecture from the origin cloud to the target cloud so that they can cease using the origin cloud. This can be done for the entirety of the customer architecture (full switch) or for parts of it where only certain existing workloads are migrated (partial switch).

---

[409] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[410] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]
[411] Ofcom / [✂] meeting, [✂], subsequently confirmed by Vodafone by email on [✂].
[412] Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂]; Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]; Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]; and Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[413] Oracle response to the MIR consultation, pages 3 and 5; and Ofcom / Oracle meeting, 15 June 2023, subsequently confirmed by Oracle by email on 9 August 2023.
[414] Ofcom / IBM meeting, 26 June 2023.
[415] The results of Public First's survey are available at: https://www.publicfirst.co.uk/files/CCIA_Survey.xlsx [accessed 19 September 2023]. The relevant questions are 49-50.

5.34 When switching between clouds, customers need to ensure their applications and data work and perform equivalent tasks in the target cloud. Therefore, technical challenges revolve around limitations to both cross-cloud interoperability and portability.

5.35 In most cases, a full switch necessitates the deployment of an integrated multi-cloud as an interim step. In addition, a partial switch can be effectively regarded as a switch from a single-cloud architecture to a multi-cloud architecture. Therefore, high technical barriers to multi-cloud will likely reinforce the technical barriers to switching between clouds.

5.36 In practice, some technical differences always exist between clouds and as a result some technical effort is always required of customers to complete the switch. However, a high degree of technical differentiation between clouds materially increases this effort as it requires customers to make additional changes to their applications and data so that they can work on the target cloud. The scale of these challenges varies depending on the use-case. Other things being equal, the main factors that determine the level of technical effort are usually the number of applications that need to be ported and the tightness of their integration into the proprietary services of the origin cloud.

5.37 The evidence indicates that technical challenges to switching are likely to be material for some customers and use-cases:

a) The qualitative part of the market research encountered few, if any, examples of organisations switching away from one of the hyperscalers.
b) The responses to our customer questionnaire confirmed the importance of technical barriers to switching, despite mitigation strategies being in place. Most customers said that technical difficulties are one of the top barriers to switching, along with skills lock-in and data egress fees.[416] For example, [✂] said that while its cloud strategy mandates

[416] [✂] said that switching cloud providers will be costly and require extensive rework, implementation and testing ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); [✂] ranked technical barriers as the second most important barrier to switching after the nature and number of apps or size of data to be ported ([✂] response dated [✂] to our customer questionnaire, question [✂]); [✂] said that switching products that it has built on one cloud provider to another is a large undertaking because of functional (capabilities of the services) and non-functional (security, resilience) differences between cloud providers ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); BT Group said that indirect switching costs (i.e. technical difficulties) often total more than any direct switching costs (i.e. data egress fees) (BT Group response to the CFI, page 9); [✂] said that, even with mitigations in place, re-engineering to an alternative cloud service would be a significantly larger endeavour in terms of the complexity, time and cost to reconfigure applications ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); [✂] ranked technical barriers as the second most important barrier to switching after the nature and number of apps or size of data to be ported ([✂] response dated [✂] to our customer questionnaire, question [✂]); [✂] ranked technical barriers as the second most important barrier to switching ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); [✂] compared switching cloud providers to switching from gas to electricity power. Despite emerging technologies to facilitate switching, it said it considers this to be 'once-in-a-decade' strategic decision ([✂]); [✂] said that limitations to or a lack of interoperability is one of the factors preventing it from switching ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); [✂] said that barriers and costs vary but the technical barriers would likely be relatively significant ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); [✂] ranked technical difficulties as the second most important barrier to switching after vendor's product performance and functionality ([✂] response dated [✂] to our customer questionnaire, question [✂]); [✂] expected lack of interoperability to be its biggest problem for migration ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); [✂] said that technical barriers to switching are significant and added only one more major obstacle - data egress fees ([] response dated [✂] to the s.174 notice dated [✂], question [✂]); Vodafone listed barriers to switching such as lack of portability standards and egress as one of the most important issues to examine in cloud infrastructure services competition (Vodafone response to the CFI, Q4.7).

the use of open standards and open-APIs, it has experienced substantial technical barriers to switching a single component of its cloud architecture[417] as moving supplier means re-architecting applications even if they are abstracted behind IaaS or CaaS (i.e. containers). This customer mentioned it followed a standard migration process, and this has been ongoing for 6 months.[418]

c) Most smaller cloud providers that engaged with our market study mentioned technical barriers to switching as one of the challenges to customer acquisition. They noted that there are limitations to interoperability and portability, which they consider materially increase switching costs and risks of lock-in.[419]

d) A [✂] report [✂] highlights a number of technical barriers to cloud portability, which appear to be in line with the technical barriers set out in the previous paragraph. These include: differences in technical capabilities, differences in processes and tools, data gravity, integration efforts and employee skills.[420]

e) The responses to our interim report and our continued stakeholder engagement has largely confirmed the importance of technical efforts as a barrier to switching between clouds.[421]

5.38 The evidence received also indicates that more mature companies and larger companies may be more affected by technical barriers to switching. This may be because these companies are more likely to have large numbers of applications and/or use various proprietary services offered by their cloud providers. In particular:

a) The market research suggests that the importance of barriers to switching may vary by company maturity. For example, 49% of the companies established for less than 2 years indicated at least one of these technical challenges as a barrier to switching: data portability, application portability and interoperability. This figure raises to an average of 58% for companies established for more than 2 years.[422] This may be because more established companies have more cloud applications.

---

[417] [✂] changed its [✂].

[418] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].

[419] BT Group response to the CFI, pages 8-10; [✂] response dated [✂] to the s.174 notice dated [✂], pages [✂] and questions [✂]; [] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].

[420] [✂]

[421] BT Group said that improving interoperability and portability between cloud providers can enable better outcomes for customers, facilitating switching, preventing lock-in (BT Group response to the interim report, page 1). [✂] said that when applications and workloads have evolved within a particular cloud provider, it becomes more difficult to move them out, which might make users unprotected against price increments, where the user might have to accept 'non-transparent' extra costs versus the cost of moving the applications and data to other platforms (Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂]). [✂] noted that, while its current architecture has been built in Kubernetes which provides it with a reasonable amount of portability, the more products that AWS release, the more options there are for a closer integration with AWS (Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂]). [✂] explained that exiting would likely be a multi-year programme which limits its options, unless it chooses to use containerisation, but given the variety of places they have workloads, it does not make sense for it to split everything they have and containerise (Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂]).

[422] Analysis of Context Consulting research data tables, Q52. After excluding users of private cloud only these figures are equal to 50% for companies established by less than 2 years and 60% on average for companies established by more than two years. More specifically, the percentage of customers indicating at least one of these technical challenges are as follows: 49% (50% excluding users of private cloud only) of companies whose

b) The market research also found some differences in how customers of different sizes perceive technical barriers to switching. The largest customers (>999 employees) were more likely to cite at least one of the following technical challenges as a barrier to switching: data portability, application portability and interoperability (64% compared to an average of 55% for companies of smaller size).[423] Larger organisations also tend to have more cloud use cases.[424] Taken together these findings indicate that larger customers were more likely to cite technical barriers to switching because they have more complex cloud architectures and more applications to port.[425]

c) Consistent with the above, the responses to our customer questionnaire suggested that technical barriers to switching are particularly strong for customers porting a large number of applications and/or a large volume of data, or porting critical infrastructure that is subject to strict regulation on resilience and security.[426] We understand that, for these customers, cloud-agnostic design often becomes simply impractical, due to the additional complexity, general time constraints and the lack of centralised coordination.

5.39 As discussed in Annex 4, to assess the share of customers that may be facing high technical barriers to switching we asked hyperscalers to provide us with the number of first-party proprietary services and of first-party PaaS services used by customers of different sizes.[427] The responses indicate that, across all hyperscalers, customers spending more than £10k per year (which represent a significant portion of the overall hyperscaler's revenues) consume at least [✂] [10-20] first-party proprietary cloud services and at least [✂] [5-15] PaaS services. We consider that customers using this number of first-party proprietary cloud services would likely face a high degree of technical complexity if they wanted to switch or replicate some or all of their architectures across multiple clouds.

5.40 The same data indicates that the number of first-party proprietary services is generally higher for customers in higher revenue bands. For example, customers spending more than $1m a year (accounting for [✂]% of hyperscalers revenue) on average take at least [✂] [30-40] first-party proprietary services and at least [✂] [20-30] PaaS services. This suggests that

business has been established by less than two years, 48% (51% excluding users of private cloud only) of companies whose business has been established by more than two but less than five years, 54% (54% excluding users of private cloud only) of companies whose business has been established by more than 5 years but less than 10 years, 59% (61% excluding users of private cloud only) of companies whose business has been established by more than 10 years but less than 20 years, 64% (68% excluding users of private cloud only) of companies whose business has been established by more than 20 years.

[423] Analysis of Context Consulting research data tables, Q52. After excluding users of private cloud only these figures are equal to 67% for companies with more than 999 employees and 56% for companies of smaller size. More specifically, the percentage of customers indicating at least one of these technical challenges are as follows: 54% for companies with 10-49 employees (59% after excluding users of private cloud only), 55% for companies with 50- 249 employees (53% after excluding users of private cloud only), 56% for companies with 250-999 employees (57% after excluding users of private cloud only), 64% for companies with more than 999 employees (67% after excluding users of private cloud only).

[424] Context Consulting research report, slide 37.

[425] For example, customers with more staff also have more complex identity and access management.

[426] See paragraph 5.37b.

[427] AWS response dated 14 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' columns E, F and D); Microsoft response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Confidential Annex 2, tab 'Q1 Customer distribution' columns E, F and D); and Google response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' columns E, F and D).

technical barriers are likely to be more material for more mature cloud users, which is consistent with the results of our market research.

5.41 Overall, we consider that customers are always likely to face some technical barriers when switching between clouds. These barriers are likely to be material for customers porting a large number of applications which are tightly integrated with proprietary cloud services and may increase further in industries that are subject to strict regulation on resilience and security. This is likely to encompass customers in many critical sectors, such as government, financial services, healthcare, social media, as well as our core sectors of broadcasting and telecoms.

### Hyperscalers services that facilitate multi-cloud and switching are not sufficient to meaningfully mitigate technical barriers

5.42 Each of the hyperscalers stated that they offer a wide range of support for multi-cloud and switching, including open-source software and standards, cloud services designed to facilitate multi-cloud and switching, and advice to customers. In particular:

a) AWS said it: (i) often builds its services on, or using, open-source technologies and standards; (ii) makes a wide range of services and tools available to customers who wish to migrate or multi-cloud; (iii) educates its customers on building for "reversibility" in their IT solutions; and (iv) allows third parties to use AWS APIs and software development kits (SDKs) outside AWS.[428]
b) Microsoft said: (i) open-source technologies (e.g. Linux, Kubernetes) facilitate switching because they are prevalent across public clouds; (ii) some commercial first-party PaaS services (e.g. Azure Arc) and some commercial third-party PaaS services (e.g. Snowflake) are designed to facilitate interoperability between clouds; (iii) it offers customers help with switching; and (iv) it makes extensive information available to developers about the services available in Azure and how customers, partners, or competitors can access that functionality.[429]
c) Google said: (i) it is committed to using open APIs across its technology stack and building many of its services on open-source solutions; (ii) it gives customers options and tools to build, migrate and deploy their applications across multiple cloud environments to avoid vendor lock-in.[430]

5.43 AWS and Microsoft said that these technologies enable customers to choose the best solution for their needs, by allowing customers to make a deliberate trade-off between portability and the value of using proprietary services when choosing their cloud architecture.[431] They offer full flexibility, and users who are generally sophisticated, make their decision in an informed way. For example:

---

[428] AWS response dated 31 October 2022 to the s.174 notice dated 24 October 2022, pages 2-3; AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, page 7, and Part A questions 7, 8, 11, 17, 23 and 31; AWS response to the interim report, paragraphs 11, 20-23; and [✂].
[429] Microsoft response to the CFI, pages 7-10 and footnote 16; Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, Part A questions 11 and 31; and Microsoft response to the interim report, paragraphs 197-203.
[430] Google response to the CFI, paragraphs 13, 14, 23 and 28; Google response dated 23 November 2022 to the s.174 notice dated 26 October 2022, Part A questions 2, 8, 17 and 31; and Google response to the interim report, paragraphs 12-13.
[431] AWS response to the interim report, paragraphs 20-22 and 27; [✂]; and Microsoft response to the interim report paragraph 180 and sections 9.2.5, 9.2.6 and 9.2.7.

a) Customers who place particular importance on minimising switching or multi-clouding costs can take basic IaaS services and run a cloud-agnostic service on top. This would facilitate multi-cloud and switching but may involve more work on the customer side and the customer may not have access to a range of innovative proprietary cloud infrastructure services offered by AWS and Microsoft.

b) Customers who particularly value innovative services offered by AWS and Microsoft may decide to build a cloud architecture using these services and accept that this may entail higher technical efforts if they decided to multi-cloud or switch. This is because, for example, an equivalent proprietary service may not be available on other clouds meaning that more reconfiguration or re-engineering efforts may be needed in case of a switch.

5.44 We agree that these technologies may in principle facilitate multi-cloud and switching. However, our evidence suggests that, while increasing, the take-up of some of these technologies is limited. In particular:

a) The number of active customers using [✂] increased from [✂] in 2019 to [✂] in 2021. These represent only around [✂]% of the number of active customer accounts using [✂] in the same years.[432]
b) The number of active customers using [✂] increased from [✂] in 2020 to [✂] in 2022. This represents only around [✂]% of the number of active customers using [✂] in the same years.[433]

5.45 We recognise that the figures on number of active customers may not precisely reflect the individual number of users.[434] However, we consider this analysis provides a good indication that take-up of some of the hyperscalers' services designed to facilitate switching and multi-cloud is likely to be limited.

5.46 Regardless of take-up, these technologies may also be insufficient on their own to facilitate multi-cloud and switching. For example, if a customer wishes to maximise interoperability and portability for its cloud architecture, it would not be enough to run all its workloads on Kubernetes (the open-source container orchestration service hyperscalers have noted) which is only one part of their cloud architecture. The customer would also need to ensure that all other parts of their cloud architecture are open-source or cloud-agnostic. In practice, as discussed above, most customers tend to use multiple cloud-specific services to build their cloud architecture. This means their cloud architectures would be equally difficult to port, whether they run on Kubernetes or not.

5.47 In addition, many cloud services that facilitate multi-cloud and switching seem to be aimed at hybrid rather than multi-cloud, meaning that they may not necessarily facilitate the take-up of more integrated forms of multi-cloud. This is because many hybrid cloud services are aimed at expanding the proprietary cloud environment to customers' on-premises IT (e.g.

[432] To carry out this analysis we have compared the total number of active customers using [✂] in 2019 and 2020 to the total number of active customers using [✂] in the same years.
[433] To carry out this analysis we have compared the total number of active customers using [✂] in 2019 and 2020 to the total number of active customers using [✂] in the same years.
[434] [✂] have explained that a number of assumptions have been required to estimate the number of active customers. For example, a unique customer might be counted multiple times due to that unique customer having multiple accounts or customer IDs. We do not consider that precise estimates of unique customers are necessary to make the point expressed above.

Amazon EKS Anywhere can be deployed on on-premises servers so that customers can effectively replicate their AWS workflows on their on-premises IT, which would facilitate migration to AWS). When applied to a multi-cloud architecture, these services may allow customers to expand their existing architecture to other cloud providers, but they would not necessarily facilitate an integrated form of multi-cloud (e.g. Amazon EKS Anywhere could be deployed on Google servers but this would not make it any easier for customers to integrate their Amazon EKS clusters with managed Kubernetes services from Google).

5.48 Moreover, the introduction of services that facilitate hybrid cloud may simply reflect cloud providers' commercial incentives to facilitate migration from on-premises legacy IT. For example, [✂].[435] As discussed below, these incentives may diminish as the market matures and the expected rate of migration from on-premises decreases.

5.49 We acknowledge that some customers may deliberately trade-off portability and business value of different cloud architectures.[436] However, to the extent technical barriers do not have a clear justification this may lead to suboptimal customer decisions.

## The evidence suggests that some of the technical barriers are not justified and might persist going forward

5.50 AWS and Microsoft acknowledged that, even with all services offered to mitigate technical barriers, a certain level of technical effort is needed when changing or adding a cloud provider. They said this is an inherent feature of IT markets rather than the result of deliberate actions aimed at locking customers in.[437] In particular, AWS and Microsoft said that:

a) Technical differentiation is the result of an active and effective competitive process:

i) AWS said that differentiation at interface or functional level is not an impediment to competition between cloud providers but rather a feature of effective competition, adding to the unique value proposition of each service and driving divergent and innovative solutions to meet a range of customer needs.[438]

ii) Microsoft said that differentiation between providers is the result of an active competitive process which had delivered many benefits to customers. To seek to limit such development (whether through mandated standardisation, required

[435] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].

[436] For example, [✂] said that it had to make a decision whether it wanted to remain highly portable or take advantage of AWS proprietary services and, given AWS are constantly releasing new innovations that [✂] competitors are using, it decided that the benefits from innovation outweighed the disbenefits of lock-in and so has not focused on portability (Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂]; and [✂] confidential response to the interim report. [✂] said it uses cloud-native technology where the benefit of doing so outweighs the cost of avoiding 'lock-in' (Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂]). [✂] said that it uses managed open-source services wherever possible because these services are more easily portable to other clouds and there is less need to update the integrations of its apps; where it consumes proprietary services it does so with open eyes, understanding the impact that this will have and the reasons for using them (Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂]). University of East Anglia said that customers, in particular of IaaS and PaaS services, are likely to account for potential portability issues during the decision-making process by, for example, testing whether applications can be ported as required (University of East Anglia response to the interim report dated 17 May 2023).

[437] AWS response to the interim report, paragraph 9; Microsoft response to the interim report, paragraphs 25 and 177.

[438] [✂].

equivalency or restrictions on open source-based services) would inevitably limit innovation and differentiation.[439]

b) Data gravity and the complexity of a multi-cloud environment may drive friction. While some mitigation techniques exist, these challenges are to some extent inevitable:

i) AWS mentioned that many of the technical challenges customers face setting up an integrated multi-cloud architecture are inherent to the integrated multi-cloud model, and are not challenges that can be ameliorated through regulatory intervention. AWS said that any data transfer between cloud providers will significantly increase latency, challenge the users' ability to track data lineage, and risk introducing data drifts. AWS said that other operational challenges inherent in integrated multi-cloud include increased security and data privacy challenges, and creating multiple points of failure for a single application.[440]

ii) Microsoft said that data-centric workloads are likely to suffer worse latency and efficiency if they are implemented with interconnections with other cloud environments. Customers can limit the impact of latency on their cloud to a certain extent by using private networks or concentrating their data and services in a single region. However, it is impossible to eliminate latency when deploying workloads over multiple clouds and/or locations as this is fundamentally an issue of physics. Microsoft added that introducing multi-cloud functionality may create single points of failure that can lower system resiliency and create complexities, including in managing procurement and costs, compliance with regulatory requirements, security.[441]

c) In addition, Microsoft said that limited interoperability between first and third-party services may in some cases be the natural result of innovation. In particular, Microsoft mentioned that cloud providers, are continuously updating their services with the latest security enhancements and innovations. An ISV or third-party cloud provider may not be able to fully test and update their applications until some time after release. As a result, there may be a reduction in interoperability between applications updated at different times, and leading to lower performance or product quality (i.e. certain features do not work correctly).[442]

5.51 We accept that some level of technical effort is always likely to be required when working with multiple clouds, multiple vendors and switching between them. Some of the technical differentiation between cloud services may be a precondition to, or the result of, cloud providers competing for customers by building new innovative products. For example, cloud providers may design their systems in the way they deem most efficient from a technical perspective, and this may differ across providers. Also, new cloud services may be inherently less interoperable or portable, if these are the result of technical innovation which is not available on all clouds. For example, if a cloud provider introduced a new cloud service based on an innovative chip that is only available to them, this service may not work as effectively (or at all) with services from other clouds.

---

[439] Microsoft response to the interim report paragraph 178.
[440] [✂].
[441] Microsoft response to the interim report paragraphs 188 and 191.
[442] Microsoft response to the interim report, paragraph 192.

5.52 However, we consider that for some of the technical barriers we have identified the justification is less clear. In particular, we have seen evidence suggesting that:

a) Differentiation between ancillary cloud services can add material complexity for customers.
b) Technical barriers are higher if first-party cloud services are not fully functional when used in combination with third-party cloud services.
c) There exists technically and commercially viable solutions to mitigate latency and integration efforts, but these are currently not available on all clouds.
d) Cloud providers may not be fully transparent about the compatibility of their cloud infrastructure services with competing open-source or ISV services.

5.53 In addition, based on the evidence examined, we consider that technical barriers might remain high going forward.

### Differentiation between ancillary cloud services can add material complexity for customers

5.54 Customers must use a set of ancillary cloud services to operationalise their cloud deployments. These will typically include billing, security, orchestration and observability services which are required irrespective of the specific cloud use-case.

5.55 We understand that many such ancillary cloud services are typically cloud-specific, meaning that customers using a cloud service from a given cloud provider must use their first party ancillary cloud services. This can materially add to the complexity of managing an integrated multi-cloud architecture since customers would be required to multiply their management efforts and continuously maintain cross-cloud consistency across billing, security, orchestration and observability. As a result, using separate sets of ancillary cloud services may decrease efficiency and increase security risks of running an integrated multi-cloud architecture (e.g. it may be less efficient to continuously monitor services across clouds, and customers may need to make sure updates to data access permissions are made consistently across user interfaces, creating scope for human error).

5.56 We note that ISVs would typically not offer a fully-fledged set of proprietary ancillary cloud services and would instead rely on those of the hosting cloud. In some cases, ISVs may offer some ancillary functionalities which are needed to take full advantage of their services (e.g. an ISV offering a database service typically would not offer a fully-fledged IAM service but may have some specific permissions/settings that they want users to be able to configure). If ISVs cannot integrate these functionalities into the ancillary cloud services of the hosting cloud, they may need to offer a separate and potentially differentiated user interface to allow customers to use those ancillary functionalities. This may increase the hassle costs of managing a multi-vendor architecture (e.g. in the example of the database vendor, customers could not control those specific permissions/settings through the cloud-specific IAM but would need to separately manage those by logging in the ISV customer profile and use the ISV-specific interface). In addition, a high degree of differentiation between ancillary cloud services of different clouds may increase the technical efforts required of ISVs to integrate with different clouds and slow down uptake of ISVs services.

5.57 In some cases, customers may want to create an integrated interface to proprietary ancillary cloud services of different vendors (e.g. to have a unified control panel for their multi-cloud architecture). This would typically be a complex (if at all possible) technical endeavour as many of the integration efforts would sit with the customers. To the extent ancillary cloud

services are differentiated, such integration efforts are likely to be higher and solutions are likely to be limited in scope.

5.58 Overall, technical differentiation between ancillary cloud services of different vendors increases the complexity of managing a multi-cloud and, to a lesser extent, multi-vendor environment. This is because customers would need to use and continuously sync multiple differentiated services or sustain material technical efforts should they wish to create a unified interface. This will in turn hinder customers' ability to switch between clouds and – to a more limited extent – within clouds.

5.59 The evidence we have received suggests that there is a high level of technical differentiation of these ancillary cloud services across clouds and across suppliers of cloud services on the same cloud. This creates additional complexity when managing a multi-cloud or multi-vendor architecture:

a) A cloud provider ([✂]) said that multi-cloud or switching can be hindered by proprietary cloud tools which are often not portable. This includes proprietary software and Day 2 operations tools[443] which are used once an application is running to e.g. monitor performance, ensure security, compliance, logging and metering. This means that an application or Day 2 operation tool has to be recreated when using another cloud provider.[444] The cloud provider mentioned that enhancing standardisation of Day 2 tools would help customers better manage their multi-cloud architectures, it would facilitate the take-up and use of ISVs services[445] and it would make it easier for ISVs to develop competing ancillary cloud services which could work across clouds.
b) Eclipse Foundation said that enabling real cloud interoperability and portability requires the implementation of a subset of cloud-neutral APIs, defined by the open-source community, on top of the existing Cloud Providers backend APIs. This would allow the open-source community to control the interoperability of ancillary cloud services between the different clouds which includes: billing, security, observability and orchestration.[446] In other words, Eclipse Foundation suggested that standardising ancillary service APIs would make it easier to set-up and manage a multi-cloud or multi-vendor environment.
c) Suppliers of professional services and customers also indicated differentiation of ancillary cloud services as one of the key drivers of the complexity associated with managing multi-cloud architectures and suggested that enhanced standardisation in this space may be beneficial:

    i) [✂] emphasised that for running an integrated multi-cloud architecture efficiently, a customer needs ancillary cloud services that provide end-to-end observability and other ancillary functions for their cloud architecture. [✂] noted that while ancillary cloud services are standard in what they do, they differ significantly on how they do it which creates added complexity for the customer. For example, each ancillary cloud service currently produces different data in a different format and with

[443] "Day 2 tools" is a term that is sometimes used in the cloud industry to refer to ancillary cloud services.
[444] [✂] confidential response to the interim report, page [✂].
[445] [✂] explained that if ancillary cloud services were standardised across first-party services and ISVs services, it would be easier for customers to create and manage architectures composed of first-party and ISVs services (Ofcom / [✂] meeting, [✂]).
[446] See section IV-B of Eclipse Foundation, 2023. Open services cloud OSC – Unlock cloud interoperability to foster the EU digital market. [accessed 4 September 2023].

different frequency. [✂] said that customers would be helped if there were more standards around ancillary cloud services. This includes standard interfaces and standards about the type, format and frequency of data produced.[447]

ii) [✂] said that IAM has been a significant barrier in taking up Google in addition to AWS. It started using AWS first and set up the AWS IAM. Several years later, when it started using Google, [✂] found it technically challenging and potentially unsafe to manage both AWS and Google IAM in an integrated fashion due to the high degree of differentiation. [✂] also had similar issues with billing services. [✂] said that avoiding the complexities associated with managing ancillary cloud services (e.g. IAM or billing) across clouds was one of the key factors driving its decision to concentrate its cloud usage on AWS.[448]

iii) [✂] noted that it is possible to integrate Okta[449] into multiple clouds, but the complexity stems from the fact that integration efforts mostly sit with the customer and entail continuous work as integrations need to be applied to any new applications/services/workloads. This means that, from an integration perspective, it may be simpler to use just one cloud but that is not necessarily the right choice.[450]

5.60 As discussed above, AWS and Microsoft submitted that technical differentiation is the result of an active and effective competitive process which drives innovation and benefits customers. In addition, Microsoft stated that seeking to limit such differentiation through mandated standardisation would inevitably limit innovation and differentiation.[451]

5.61 Overall, we consider that the degree of differentiation between ancillary cloud services can materially contribute to the technical efforts required of customers when implementing more integrated forms of multi-cloud, and – to a lesser extent –multi-vendor architectures. This will in turn hinder customers' ability to switch between clouds and – to a more limited extent – within clouds.

5.62 We recognise that technical differentiation between ancillary cloud services may have not been designed with the aim to increase technical barriers and may simply reflect different approaches or preferences across providers. At the same time, the current level of differentiation between ancillary cloud services can add material complexity to the management of multi-cloud and multi-vendor architectures (and switching as a result).

5.63 We understand that many ancillary cloud services largely exhibit the same functionalities. This means that less technical differentiation between these ancillary cloud services would unlikely limit customer access to innovative ancillary cloud functionalities and would unlikely restrict cloud providers' and ISVs' ability to innovate over the core functionalities of their services. In line with this, evidence received indicates there would be some industry and customer support for greater harmonisation efforts in this space.

---

[447] Ofcom / [✂] meeting, [✂].
[448] Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂].
[449] Okta is an ISV which supplies cloud services, including some ancillary cloud functionalities.
[450] Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂].
[451] AWS response to the interim report, paragraph 24. Microsoft response to the interim report, paragraph 178. We discuss these submissions in more detail and respond to them at paragraphs 5.50 and 5.51 respectively.

## Technical barriers are higher where first-party cloud services are not fully functional when used in combination with third-party cloud services

5.64 We have also seen evidence that AWS and Microsoft make some functionalities of their cloud infrastructure services available only to their own cloud infrastructure services and not to competitors. In other words, in some cases the full set of functionalities of a certain service from AWS or Microsoft is only available when this service is used in combination with another first-party service. We refer to this type of restriction as 'asymmetry of functionalities'.

5.65 We understand that, in practice, AWS and Microsoft impose such restrictions by designing their cloud infrastructure services and interfaces so that it is not possible to take advantage of certain functionalities when using them in combination with a third-party cloud service. This would seem to particularly affect customers' ability to import, export and exchange data between first- and third-party cloud services, because of limitations to the available sources and destinations on API level. For example, a customer wanting to use a certain functionality of an analytics cloud service (e.g. upload data) would typically do so by making a request to the API of the analytics service (e.g. upload 'data.file' from 'URL'). However, if the URL can only refer to a location within the same cloud, customers would not be able to upload data from a source located on a different cloud.

5.66 When this is the case, customers are prevented from building or switching to an integrated multi-cloud or multi-vendor architecture involving such functionalities. For example, if an analytics service offered by cloud A can only load data from a first-party storage service offered by the same cloud, customers[452] will not be able to set-up or switch to an integrated multi-cloud architecture combining the analytics service offered by cloud A with the storage service offered by cloud B.[453]

5.67 An asymmetry of functionalities may be temporary when the full set of functionalities is initially restricted to first-party cloud services and later made available to third-party cloud services from other clouds or ISVs.[454] As discussed below, in such cases customers may still be discouraged to create a multi-cloud or multi-vendor architecture if they expect a delay in accessing the latest functionalities introduced by the hosting cloud.[455] As such, even a temporary asymmetry of functionalities may discourage take-up of multi-cloud and multi-vendor architectures.

5.68 As discussed above, we accept that the ability to add and monetise proprietary functionalities can drive innovation and benefit customers. However, we do not consider this ability is necessarily put at risk where customers take advantage of these functionalities

---

[452] This would include customers taking up storage and analytics from cloud A who want to switch to the storage solution of a different provider, as well as, customers taking up storage and analytics from cloud B who want to switch to the analytics solution of cloud A.

[453] The only way a customer may be able to use the analytics service from cloud A in combination with the storage service from cloud B would be to also purchase storage from cloud A and copy data into that any time data needed to be analysed. However, this would lead the customer to purchase and pay for two storage solutions and the resulting multi-cloud set-up would be less integrated. In fact, we note this latter scenario may be more akin to a siloed multi-cloud which, as discussed above, some customers may be implementing due to the challenges of the integrated multi-cloud.

[454] For example, as discussed below and further detailed in Annex 4, we understand that the limits to interoperability that we had identified in the interim report in relation to Amazon SageMaker were recently lifted.

[455] See paragraphs 5.70a and 5.76.

in combination with a third-party cloud service. For example, a provider designing a new 'upload' functionality for its analytics service could monetise such addition (e.g. by increasing the price of the service or by charging customers using the new functionality) irrespective of whether data are uploaded from a first-party or a third-party source.

5.69 Based on an analysis of publicly available documentation, we have identified a number of AWS and Microsoft services that appear to be subject to such restrictions. These services are listed below along with quantitative or qualitative information on the uptake of these services.

**Table 5.2: Examples of AWS and Microsoft cloud services with potential interoperability limits**

| Supplier | Service | Uptake | Potential limits to interoperability[456] |
|---|---|---|---|
| **AWS** | Amazon Athena | [✂].[457] | Can only query data stored on Amazon S3. |
| **AWS** | Amazon IVS | This service is used by large companies such as GoPro.[458] | Can only auto-record to Amazon S3. |
| **AWS** | Amazon Kinesis Video Streams | [✂].[459] | Can only deliver extracted images to Amazon S3. |
| **AWS** | Amazon Omics | AWS website shows this service is popular with customers in the clinical space.[460] | Uses Amazon S3 for data import and export. |
| **AWS** | Amazon Pinpont | [✂].[461] | Allows adding Amazon Personalize recommendations to a marketing email campaigns, but not from third-party recommendations engines. |
| **AWS** | Amazon RedShift | $[✂] ([✂]% YoY growth)[462] [✂].[463] | Can only bulk load data from Amazon S3. |

[456] For AWS see: Amazon Athena; Amazon Interactive Video Service; Amazon Kinesis Video Streams; Amazon Omics; Amazon Pinpoint; Amazon RedShift; Amazon SageMaker; Amazon SageMaker DataWrangler; Amazon Timestream; and AWS IoT Events Documentation [accessed 28 September 2023]. For Microsoft see: Azure Stream Analytics; and Azure IoT Hub [accessed 28 September 2023].
[457] AWS response dated 28 November2022 to the s.174 notice dated 24 October 2022, Part A question 33, Annex A Q33.1.11, p.84.
[458] AWS website. Interactive Live Streams – Amazon Interactive Video Service [accessed 28 September 2023].
[459] AWS response dated 28 November2022 to the s.174 notice dated 24 October 2022, Part A question 33, Annex A Q33.1.11, p.84.
[460] AWS website. Genomic Data Analysis – Amazon Omics Customers [accessed 28 September 2023].
[461] AWS response dated 28 November 2022 to the s.174 notice dated 24 October 2022, Part A question 33, Annex A Q33.1.9, p.31.
[462] [✂]. AWS response dated 13 January 2023 to question 3 of our follow-up email dated 16 December 2022 concerning the s.174 notice dated 24 October 2022, Part B question 4 (Annex Q2.2).
[463] AWS response dated 28 November2022 to the s.174 notice dated 24 October 2022, Part A question 33 (Annex A Q33.1.11, p.84).

| Supplier | Service | Uptake | Potential limits to interoperability[456] |
|---|---|---|---|
| **AWS** | Amazon SageMaker [464] | $[✂] ([✂]% YoY growth)[465] | Can only access training data from Amazon S3, Amazon EFS and Amazon FSx. |
| **AWS** | Amazon SageMaker Data Wrangler | This product is ancillary to Amazon SageMaker. | Can only import data from Amazon S3, Amazon Athena, Amazon Redshift, Snowflake, and Databricks. |
| **AWS** | Amazon Timestream | $[✂] ([✂]% YoY growth)[466] | Can only use AWS Backup service to manage backups. |
| **AWS** | AWS IoT Events | $[✂] ([✂]% YoY growth)[467] | Can only trigger actions with other AWS services. |
| **Microsoft** | Azure Stream Analytics | $[✂] ([✂]% YoY growth)[468] | Exclusively support native Azure services as inputs. |
| **Microsoft** | IoT Hub | $[✂] ([✂]% YoY growth)[469] | Allows basic interoperability but prioritises integration with other Azure service. |

*Source: Ofcom.*

5.70 We also heard concerns from some cloud providers that asymmetry of functionalities may hinder customers' ability to multi-cloud and switch:

a) A cloud provider ([✂]) said although innovation increases the quality of service for customers, it may also increase technical switching barriers if improvements are only available to first-party cloud services of a cloud provider with market power, but not to competing third-party cloud services. The cloud provider added that there needs therefore to be ongoing and timely symmetry of the functionalities available to first and third-party cloud services, in order to ensure competition and true customer choice in a multi cloud environment. [470]

b) Another cloud provider ([✂]) discussed an example of an AWS cloud infrastructure service (SageMaker) which is more functional on AWS cloud. It mentioned that to realise the benefits of AWS SageMaker, it is necessary to consume the wider AWS ecosystem of

[464] AWS has lifted the restrictions on this cloud service that we identified at the interim report stage. See paragraph 5.72 and Annex 4.
[465] Ofcom analysis of AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, Part B question 4, Annex 2.
[466] AWS response dated 28 November2022 to the s.174 notice dated 24 October 2022, Part A question 33, Annex A Q33.1.26, p.38.
[467] Ofcom analysis of AWS response dated 12 December 2022 to the s.174 notice dated 24 October 2022, Part B question 4, Annex 2. [✂].
[468] Ofcom analysis of Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, Part B question 4, Annex B1-4.
[469] Ofcom analysis of Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, Part B question 4, Annex B1-4.
[470] [✂] confidential response to the interim report, page [✂]; Ofcom / [✂] meeting, [✂].

services, leading to lock-in, and discouraging users from selecting best-of-breed technologies across the market. In addition, ([✂]) raised concerns that AWS may not be fully transparent about what this means so that customers may not be fully aware of such restrictions when they first take-up SageMaker. [471]

c) In their report on the cloud market, the Dutch competition authority (ACM) stated that larger cloud providers may keep software standards of their services closed in order to limit the interoperability with services of other providers. In this context, the ACM referred to a meeting with VMware suggesting that, as a result of these limitations, connecting third-party services with cloud-native services is not always possible.[472]

5.71 AWS and Microsoft disagreed with our analysis. We summarise their arguments and our responses below. We set these out more comprehensively in Annex 4.

*AWS response*

5.72 In response to the interim report, AWS noted that the interim report only identifies interoperability limitations in a small subset of its cloud services, focusing on 10 services for which there are competing software solutions available that customers can run on AWS (or elsewhere). In addition, AWS said that the features identified as limiting interoperability are described inaccurately, exist alongside features that ensure interoperability, or are the product of an objective technical limitation.[473]

5.73 We disagree with AWS that the features identified as limiting interoperability are described inaccurately. Having reviewed AWS submission in detail, we consider that AWS has not provided any evidence to rebut our analysis. The only exception to this is Amazon SageMaker, where the limits to interoperability that we identified in the interim report were lifted after its publication. As discussed above, while we welcome the lifting of such restriction, we consider customers may still be discouraged to create a multi-cloud or multi-vendor architecture if they anticipate that the full set of functionalities would only be usable within their architecture after some time.

5.74 We also disagree with AWS that the limitations identified are not important since they exist alongside other features that ensure interoperability. We consider that limitations to effective interoperability may still exist, if the features that enable interoperability are less efficient compared to the feature being restricted to the first-party cloud service. For instance, in relation to Amazon RedShift, while other data ingestion options may exist, none of them are comparable to the speed and efficiency of bulk-loading data (which is limited to Amazon S3 only). Bulk-upload is also a very commonly used feature: large data sets are routinely loaded into data warehouses. Those processes can be essential for time critical applications.

5.75 We also disagree with AWS that the limitations identified are just the product of an objective technical limitation. Our analysis suggests that it is not technically difficult to provide access to all functionalities identified in the table above. More generally our analysis indicates that, for a given cloud service, it is in most cases technically possible to ensure that the full set of functionalities that is available to other first-party cloud services can also be made available to third-party cloud services. This outcome may even be easier to accomplish

[471] Ofcom / [✂] meeting, [✂].
[472] The Dutch competition authority (ACM), May 2022. Market Study Cloud Services, page 61. [accessed 28 September 2023].
[473] AWS response to the interim report, paragraphs 24-26.

if all functionalities were made accessible from the design stage (i.e. 'interoperability by design').

*Microsoft response*

5.76 Microsoft's response to the interim report recognised that temporary limitations of interoperability between first and third-party cloud services can discourage take-up of multi-vendor or multi-cloud architectures. However, Microsoft attributes these limitations to the natural result of innovation from cloud providers which may take some time to be incorporated and enabled by ISVs or other cloud providers.

5.77 In addition, Microsoft said that it does not have a strategy to frustrate interoperability and that it is not accurate to characterise the services identified in the interim report as not interoperable. In relation to Azure Stream Analytics, Microsoft said it configured the data source to be Azure services because Azure Stream Analytics is designed to analyse and process large volumes of data with sub-millisecond latencies. Microsoft also said that, despite this, Azure Stream Analytics is still extendable to other clouds through connectors (e.g. by configuring one of the Azure data sources, Even Hub, to take data from other sources).[474] In relation to IoT Hub, Microsoft said that, once the data is in Azure, the customers can choose to send that data across to another cloud for processing if desired. [475]

5.78 Some limitations to full interoperability may be the natural by-product of innovation. However, we are concerned that in some cases – such as in the case of Azure Stream Analytics and IoT Hub – these limitations are instead the result of a deliberate service or interface design which would not allow, even if just temporarily, customers to take advantage of certain functionalities when combining first- and third-party cloud services.

5.79 We understand why Microsoft would want to ensure the low latency feature of Azure Stream Analytics, but we do not consider this argument to be sufficient for limiting the data sources to Azure services. First, we note that Google Cloud Dataflow (Google's competing product) is allowing customers to define any data source, which suggests that there may be technically feasible ways to do this notwithstanding latency. Second, the latency argument does not explain why customers are not able to define the data source to be a third-party cloud service hosted on Azure. This could be particularly beneficial for customers wishing to create a more cloud-agnostic architecture. Finally, Microsoft's view is premised on latency being an inherent feature of the market. But as noted below, there are ways in which cross-cloud latency could be mitigated.

5.80 We also disagree with Microsoft that these cloud services are still expandable. As noted in relation to AWS above, we consider that limitation to effective interoperability may still exist, if the other features available to customers are less efficient compared to the one being restricted to the first-party cloud service.

5.81 In a follow-up meeting with Ofcom, Microsoft also said that in some cases limitations on interoperability are due to the fact that two services are really part of the same service which was split in two to allow easier "plumbing" (i.e. more granular updates to avoid big crashes of customer systems).[476] As discussed above, we accept that cloud providers should be allowed to design their systems however they deem most efficient, which may include adopting a less modular approach to the design of their services. However, in this case we

[474] Microsoft response to the interim report paragraphs 212 and 213.
[475] Microsoft response to the interim report, paragraphs 214-216.
[476] Ofcom / Microsoft meeting, 15 June 2023.

note that Microsoft markets Azure Stream Analytics as a separate cloud service and that other cloud providers and ISVs offer comparable analytics services on a standalone basis. This would suggest that it is technically possible to separate out the two services and there is likely to be demand for the separate parts.

5.82 Overall, we consider that technical barriers are higher where first-party cloud services are not fully functional when used in combination with third-party cloud services (i.e. asymmetry of functionalities). Depending on whether it applies to any third-party cloud service or only to third-party cloud services hosted on a different cloud, asymmetry of functionalities can discourage both integrated multi-cloud and multi-vendor architectures. Asymmetry of functionalities would not appear to be a necessary precondition to or the natural result of healthy competitive dynamics.

### There may be technically viable solutions to mitigate data gravity and integration efforts but these are not currently available on all clouds

5.83 As discussed above, AWS and Microsoft recognised that some customers may face technical challenges when using multiple clouds or switching. However, they told us that many of these challenges (e.g. data gravity, integration efforts) are inherent traits of cloud technologies, suggesting that it is not technically possible to mitigate them.

5.84 We are aware of examples that suggest, in principle, there are technically and commercially viable solutions to mitigate some of the technical barriers stemming from data gravity and integration efforts. In particular, we understand that connecting data centres of different cloud providers is likely to be an effective solution to mitigate cross-cloud data gravity and multi-cloud integration efforts:

a) Oracle signed a partnership with Azure which involves directly connecting their data centres and integrating their virtual environments so that customers can seamlessly use services from Azure or Oracle.[477] We understand that this solution makes it easier for customers to mix-and-match services from Oracle and Microsoft. Connecting their data centres mitigates data gravity issues, whereas the integration of Oracle and Microsoft virtual environments mitigates integration efforts on the customer side.
b) Most cloud providers, including AWS and Microsoft, offer the possibility for customers to directly connect on-premises architectures with their clouds to mitigate latency and other data gravity issues (e.g. data security or governance).[478] We also understand that there may be ways for customers to indirectly interconnect multiple clouds. For example, customers could connect AWS to their own data centre (or to a data centre provided by a third-party) and then connect this intermediary datacentre to e.g. Microsoft.[479] We note however that adopting such solutions would be very technically complex and potentially less efficient than a direct interconnection between clouds.

[477] See Microsoft website. Oracle and Microsoft announce availability of Oracle Database Service for Microsoft Azure - Stories [accessed 21 September 2023].
[478] For example, [✂] explained that it has an integrated hybrid environment with AWS and that ensuring there is no lag between these can be challenging but there are services offered by AWS (such as, the direct connect) that help mitigate these challenges (Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂]).
[479] For example, we understand that Microsoft provides ExpressRoute, a private connection between Azure datacentres and a customer's network or another cloud provider's equivalent private connection. See Microsoft response to the interim report, paragraph 189. See also Microsoft website. Connectivity to other cloud providers [accessed 14 September 2023].

c) We understand that most cloud providers, including AWS and Microsoft, typically have multiple datacentres in each availability zone. These are located in separate geographic areas but are physically interconnected to ensure greater resiliency (e.g. if one datacentre has an issue, application and workloads can be quickly ported to a separate datacentre in the same availability zone), lower latency and more capacity (e.g. AWS mentioned that in some cases a customer may be able to run different parts of an application in different closely-located datacentres of the same availability zone).[480] This further corroborates the view that physical interconnection of datacentres within geographical regions is a technically feasible way to mitigate data gravity issues.

5.85 In addition, we observe that data gravity and integration efforts may also be mitigated by making first-party cloud infrastructure services available on other clouds. For example, IBM is making their cloud infrastructure services available as cloud-native services in any cloud data centres, including AWS and Azure.[481] Google is also making some of their services available on other clouds (e.g. BigQuery Omni).[482] In such cases, IBM and Google (limited to BigQuery Omni) are effectively acting as ISVs on other clouds. From a customer's perspective this means that implementing a cloud architecture involving a mix of such services becomes more akin to adopting a multi-vendor architecture. As set out above, this is technically easier since services hosted on the same cloud are easier to integrate and are not subject to data gravity issues (e.g. latency). In addition, making services available across clouds would facilitate switching between clouds as customers could more easily replicate parts of their architecture in the target cloud.

5.86 Overall, some of the key technical challenges faced by customers (namely, integration efforts and data gravity) could be mitigated through increased industry adoption of some of the above-mentioned solutions which appear to be technically and commercially viable. However, such solutions are not widely available across clouds – and particularly between AWS and Azure – which can add to the barriers to more integrated forms of multi-cloud and switching.

### Cloud providers may not be fully transparent about the compatibility of their cloud infrastructure services with competing services

5.87 We have received submissions suggesting that all hyperscalers, but particularly AWS and Microsoft, may be exacerbating technical barriers by not being fully transparent about the extent to which their cloud infrastructure services are compatible with competing ISVs' or open-source services. As a result, customers using these services may not be aware of the extent to which they may need to rewrite some of their code if they wish to multi-cloud or switch.[483] More specifically:

a) An ISV ([✂]) explained that AWS and Microsoft have developed proprietary cloud services (namely [✂]) which they sell as compatible to its cloud service (namely [✂]). Customers may take-up these first-party cloud services because they are more visible, easy to use, or because customers are under the impression – driven by a lack of

[480] Ofcom / AWS meeting, 25 July 2023.
[481] Ofcom / IBM meeting, 26 June 2023.
[482] Google response dated [✂] to the s.174 notice dated [✂], Part [✂], question [✂], page [✂].
[483] As noted in Section 7, while AWS and Microsoft may offer both the first- and third-party cloud infrastructure services, customers may be induced to take-up the first-party ones because AWS and Microsoft make them more prominent, integrated, easy to use and/or because they may not fully disclose the degree of differentiation of their proprietary software.

transparency from the hyperscalers – that these cloud services are fully compatible with (or even a version of) the ISV's cloud service. However, according to this ISV, the compatibility of these first-party cloud services is limited, meaning that a customer using them may need to considerably reconfigure their application if they wish to switch to the ISV's service.[484]

b) Another ISV ([✂]) explained that the hyperscalers build their services on open standards and open APIs but "tweak"[485] them to make it more difficult for customers to switch as they would need to rewrite much of their code. They explained that this kind of action is presented by the hyperscalers as if they were in favour of open standards, but in fact constitute a barrier to switching and moving data.[486]

c) A cloud provider ([✂]) noted that certain market players such as Microsoft have taken steps that limit interoperability or compatibility with other cloud providers, such as by modifying open-source software to include their own, non-open APIs, or building upon open standards in a way that creates lock-in. [✂] said these actions make it harder for customers and ISVs to switch to a broadly equivalent service from another provider without first having to significantly change how they deploy their service.[487]

d) Another cloud provider ([✂]) explained that hyperscalers could have incentives to market a solution as more open and interoperable than it is in reality to capture significant market share before progressively restraining access to its technology. In this context, it gave the example of TensorFlow, a service which is marketed as open by Google.[488]

e) Another cloud provider ([✂]) said that there is an issue with some cloud providers marketing their clouds in a non-transparent way. It said that such cloud providers say they offer a lot of open-source software which is actually closed-source.[489] Where this is the case, customers may be led to believe they are using open-source software when indeed some of the features are proprietary, which may limit their ability to switch and multi-cloud down the line.

5.88 AWS stated that it does not tweak open-source software but offers managed services for such software which enables customers to limit the technical effort in running and maintaining it. AWS said that, while managed open-source services include plug-ins that integrate with other AWS services, changes to the underlying open source are documented and do not prevent customers from easily transitioning to their next solution.[490] AWS also said it aims to keep open-source software in managed services as close to the upstream project versions as possible.[491]

5.89 Microsoft stated that differentiation between providers is the result of an active competitive process including through services based on open-source. Mandating restrictions on open source-based services would inevitably limit innovation and differentiation.[492]

---

[484] Ofcom / [✂] meeting, [✂].
[485] We understand that hyperscalers do not "tweak" open standards and open APIs but instead recreate them with some functional differences.
[486] [✂]; and Ofcom / [✂] meeting, [✂].
[487] [✂].
[488] [✂] response dated [✂] to s.174 notice dated [✂], Part A [✂]; [✂] confidential response to the interim report, pages [✂]; and Ofcom / [✂] meeting, [✂].
[489] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[490] AWS response to the interim report, paragraph 23.
[491] [✂].
[492] Microsoft response to the interim report paragraph 178.

5.90 We acknowledge that open-source licences allow modifications of the software. We also consider that such modifications can be beneficial to customers, for example because they introduce innovative new features. Therefore, we are not concerned with the development of managed open-source services. However, we consider a lack of transparency around such modifications can increase technical efforts faced by customers wanting to build a multi-cloud or multi-vendor architecture and switch.

5.91 We recognise that hyperscalers typically publish documentation around the compatibility of their first-party proprietary services with competing ISVs' or open-source services. However, we have also seen evidence suggesting that in some cases the degree of compatibility may not be unequivocally clear from such documentation. For example, AWS has a developer guide on Amazon DocumentDB stating that "the compatibility of this service with MongoDB 4.0 means that a vast majority of the applications, drivers, and tools customers already use today with their MongoDB 4.0 databases can be used with Amazon DocumentDB 4.0 with little or no change".[493] MongoDB contests this statement arguing that "compatibility testing reveals it fails over 64% of the MongoDB API correctness tests. Applications written for MongoDB will need to be re-written to work with Amazon DocumentDB".[494]

5.92 Overall, we remain concerned that hyperscalers may not always be fully transparent about the compatibility of their cloud infrastructure services with competing ISVs' or open-source software. Lack of such transparency may add to the technical challenges faced by customers who may not be aware of the extent to which they may need to rewrite their code if they wish to build a multi-cloud or multi-vendor architecture and switch.

### Technical barriers might remain high going forward

5.93 AWS and Microsoft said that there are several reasons why they have an incentive to be fully interoperable and lower tech barriers:

a) AWS and Microsoft acknowledged that more interoperability would lower the tech efforts needed to migrate on-premises workloads to the cloud.[495]
b) Microsoft said that offering full interoperability allows customers using other clouds to connect to and use Microsoft services. Hence, restricting interoperability would not be profitable as it would prevent Microsoft from winning these workloads.[496]
c) AWS mentioned that customers demand the flexibility to incorporate third-party technology and services into their cloud solutions. If customers are not confident that they will be able to do so when the next innovation is released, or when they need a niche solution, they will not choose AWS.[497] Similarly, Microsoft stated that making APIs available for developers to build on top of ultimately enhances the offer to Azure customers.[498]

5.94 All cloud providers have an incentive to interoperate with customers' legacy on-premises IT to increase demand for their clouds. However, we consider these incentives may not be sufficient to ensure full interoperability across clouds and we are concerned that such

[493] See AWS website. MongoDB Compatibility - Amazon DocumentDB [accessed 14 September 2023].
[494] See MongoDB website. Comparing Amazon DocumentDB And MongoDB [accessed 14 September 2023].
[495] [✂]; and Microsoft response to the interim report, paragraph 181.
[496] Microsoft response to the interim report, paragraphs 181 and 209.
[497] [✂].
[498] Microsoft response to the interim report, paragraphs 181 and 209.

incentives may diminish as the market matures and the expected rate of migration from on-premises decreases.

5.95 We also recognise that to ensure usability of their clouds and/or in response to customer demand, all cloud providers, including market leaders, have had to grant a certain level of openness and interoperability which is beneficial for customers and competition. However, the evidence examined above suggests that in practice, incentives to interoperate with other clouds may be relatively weaker compared to incentives to interoperate with on-premises legacy IT (e.g. direct connect is typically offered for hybrid but not for multi-cloud). In addition, compared to smaller clouds, AWS and Microsoft may have a weaker incentive to respond to customer demand for cross-cloud interoperability. This may be because compared to smaller clouds, AWS and Microsoft may have a stronger incentive to defend their installed base and to leverage other advantages (such as product range, scale, and network effects) to attract new customers into their ecosystems.

5.96 In addition, the evidence examined indicates that the nature of this market is such that, where profitable, cloud providers, and particularly AWS and Microsoft, may be able to resist market forces pushing for lower technical barriers. First, interoperability is a relational concept (i.e. it takes at least two to be interoperable). This means that AWS and Microsoft (as any other cloud provider) have the ability to unilaterally restrict interoperability levels of their competitors (e.g. if AWS decides that Amazon RedShift cannot bulk-upload from Google Cloud Storage there is nothing that Google can do). Second, standardisation efforts would need a critical mass to succeed. Given their market position in cloud, both AWS and Microsoft would likely have the ability to slow-down or block any standardisation efforts should they decide to not fully support it.

5.97 Overall, this suggests that despite customer demand for interoperability some technical barriers might remain high, thereby restricting customers' ability to multi-cloud and switch.

## Conclusion on technical barriers

5.98 Based on the evidence received, we consider that there are several factors which can reduce interoperability and portability in cloud and, as a result, hinder customers' ability to implement different multi-cloud or multi-vendor architectures and switch. These include: technical differentiation between cloud infrastructure services, technical differentiation between ancillary cloud services, integration and operationalisation efforts, asymmetry of functionalities, data gravity, lack of technical skills and lack of transparency.

5.99 In practice, the overall scale of technical efforts required of customers will typically be a result of a combination of the above challenges, which is likely to be particularly significant for customers wishing to integrate services from multiple clouds (i.e. integrated multi-cloud) or switch between them.

5.100 Some customers may be relatively unaffected by some of these factors. This includes those who use a few applications built on common or sufficiently similar standards. However, the technical efforts are likely a material barrier to switching and multi-cloud for customers with large numbers of applications or with a cloud architecture that is tightly integrated with many first-party cloud services. Based on the results of the market research and on cloud providers' data, we believe this is likely to encompass a material portion of customers across several sectors and reflect a significant fraction of revenue.

5.101 Each of the hyperscalers offer cloud services which can mitigate technical barriers to switching and multi-cloud, including the use of open-source software and open standards which are available across clouds. However, our evidence indicates that the nature and uptake of these services is generally not sufficient to meaningfully offset technical barriers to multi-cloud and switching.

5.102 We accept that some technical barriers may be a precondition to, or the natural result of cloud providers seeking to differentiate themselves by innovating to the benefit of customers. However, we consider that for some of the technical barriers we have identified the justification is less clear.

5.103 In particular, our evidence indicates that: (i) differentiation between ancillary cloud services can add material complexity for customers; (ii) technical barriers are higher if first-party cloud services are not fully functional when used in combination with third-party cloud services; (iii) there may be technically and commercially viable solutions to mitigate data gravity and integration efforts but these are not currently available on all clouds; and (iv) cloud providers may not always be fully transparent about the compatibility of their cloud infrastructure services with competing open-source or ISVs' services.

5.104 Some of these technical challenges can also add to the barriers to entry and expansion faced by ISVs and as a result further limit customers' ability to multi-cloud and switch. Some services offered by ISVs can mitigate technical barriers by facilitating multi-cloud and by allowing customers to build cloud-agnostic architectures which are easier to port. If ISVs find it more difficult to enter or expand, this could deprive customers of a potential solution for building more interoperable or portable cloud architectures.

5.105 Overall, these technical barriers can hinder customers' ability to work with multiple clouds, switch between them and – to a lesser extent – implement multi-vendor architectures within the same cloud. Moreover, the evidence examined would suggest that despite demand from customers and a push from some market participants to lower technical barriers, cloud providers – and especially market leaders – may be able to keep technical barriers high in the foreseeable future.

# Data egress fees

5.106 In this subsection we examine the extent to which data egress fees may hinder customers' ability to switch and multi-cloud. The subsection is structured as follows:

a) We define the term egress fees, and set out factual background on the current level of egress fees, the current aggregate volume of data transfers and associated provider revenue;
b) We explain the evidence that egress fees may hinder customers' ability to multi-cloud and how this may differ between different customer groups;
c) We explain the evidence that egress fees may hinder customers' ability to switch between cloud providers and how this may differ for different types of customers; and
d) We consider indicators as to the relationship between the egress fees which are currently charged by hyperscalers and the costs of providing egress.

## What are egress fees?

5.107 Customers may need to transfer their data:

a) into a cloud provider's infrastructure ("ingress");
b) within a cloud provider's infrastructure ("internal data transfer"); and
c) out of a cloud provider's infrastructure ("egress").

5.108 Customers may egress data out of a provider's cloud for a variety of reasons, including:

a) to deliver content to end users or applications;
b) to move data between the cloud and its on-premises data centres; and
c) to move data between different cloud providers, either as part of a multi-cloud architecture, or as part of switching between two cloud providers.

5.109 Some providers charge customers when they transfer data out of their cloud. These charges apply irrespective of the purpose of the data transfer out. We refer to these charges as egress fees.

## Egress usage and charges

5.110 Table 5.3 below sets out the current volume of data transferred internally within each hyperscaler's cloud and externally out of each hyperscaler's cloud by UK customers in 2022. It also shows the revenue that each hyperscaler received in relation to these data transfers, and how this relates to total cloud revenue. In 2022, egress fees accounted for a relatively small proportion of hyperscalers' UK public cloud revenues. Hyperscalers differ as to whether their UK customers transferred more data within their clouds or out of their clouds. However, all three received more revenue from external data transfers than from internal data transfers from their UK customers, partly because internal data transfers are charged at a lower price.

**Table 5.3: Data transfer volumes and revenues from UK customers, 2022**[499]

| | AWS | GCP | Microsoft Azure |
|---|---|---|---|
| **Data transferred internally within the cloud** | [✂] GB | [✂] GB | [✂] GB |
| **Data transferred out of the cloud** | [✂] GB | [✂] GB | [✂] GB |
| **Internal data transfer fees received** | £[✂] | £[✂] | £[✂] |
| **Egress fees received for data transfers out** | £[✂] | £[✂] | £[✂] |

[499] Sources: Ofcom analysis based on responses to s.174 notices. AWS response dated 1 August 2023 to s.174 notice dated 13 June 2023, questions 1i and 1j; AWS response dated 31 March 2023 to s.174 notice dated 24 October 2022, Part B question 4. Google response dated 11 July 2023 to s.174 notice dated 13 June 2023, questions 1i and 1j; Google response dated 31 March 2023 to s.174 notice dated 26 October 2022, Part B questions 4 and 7. Microsoft response dated 9 December 2022 to our s.174 notice dated 21 October 2022, Part B questions 4, 6 and 7; and Microsoft response dated 25 April 2023 to our s.174 notice dated 21 October 2022, Part B questions 4, 6 and 7. Fees received converted to GBP based on average exchange rate Jan-Dec 2022.

| | | | |
|---|---|---|---|
| **Egress fee revenue as a proportion of total UK public cloud revenue** | [✂]% | [✂]% | [✂]% |

*Notes: 1. For one hyperscaler [✂] the precise definition of UK customers varies across the datapoints shown in this table. Due to limitations in the available data, the estimates for data transfer volumes and data transfer revenues use a narrower definition of UK customers, whereas the datapoint for egress fee revenue as a proportion of total UK public cloud revenue uses a wider definition of UK customers (in both the numerator and denominator). 2. One hyperscaler [✂] told us it meters and charges some of its data transfer volumes using units other than gigabytes, and that it cannot straightforwardly capture these volumes in gigabyte units. The total data transfer volumes presented for that hyperscaler consequently capture gigabyte volumes only, and therefore understate total data transfer volumes. 3. One hyperscaler [✂] provided data transfer volumes in gibibytes (GiBy), and we have converted its provided estimates into gigabytes (GB) to allow for comparison with the other hyperscalers. 4. One hyperscaler [✂] was not able to provide a fully comprehensive breakdown of internal and external data transfer volumes, so the volumes shown above are likely to understate its total data transfer volumes. There are also some differences between hyperscalers regarding the inclusion or exclusion of egress volume and egress revenue involving data transfers made via private network connections. However, taking wider account of methodological similarities and differences, we consider that the datapoints shown are broadly comparable across the hyperscalers.*

5.111 Figure 5.4 below shows the list prices for egress charged by six cloud providers for transferring data out of the cloud via the public internet. We understand that both Google and Microsoft charge higher prices than those shown for egressing data through their own networks.[500]

5.112 The figure illustrates that list prices for the hyperscalers and IBM may include:

a) a "free tier" whereby customers are able to egress a certain volume of data without paying any fees (although some providers do not offer a free tier);
b) a higher marginal price for low volumes of data transfer; and
c) declining marginal prices for higher volumes of data transfer.

5.113 Figure 5.4 also shows that the hyperscalers and IBM charge egress fees which are around 5-10 times higher than OVHcloud and Oracle. Both Oracle and OVHcloud also have a flatter fee structure, with only one tier of egress fees with non-zero prices.

[500] See Microsoft website, Routing preference in Azure - Azure Virtual Network | Microsoft Learn, Network Service Tiers overview|Google Cloud and Google website, Google Cloud networking in depth: Understanding Network Service Tiers | Google Cloud Blog [accessed 1 September 2023].

**Figure 5.4: Egress fee list prices for transferring data out via the public internet, September 2023 (GBP per GB)**

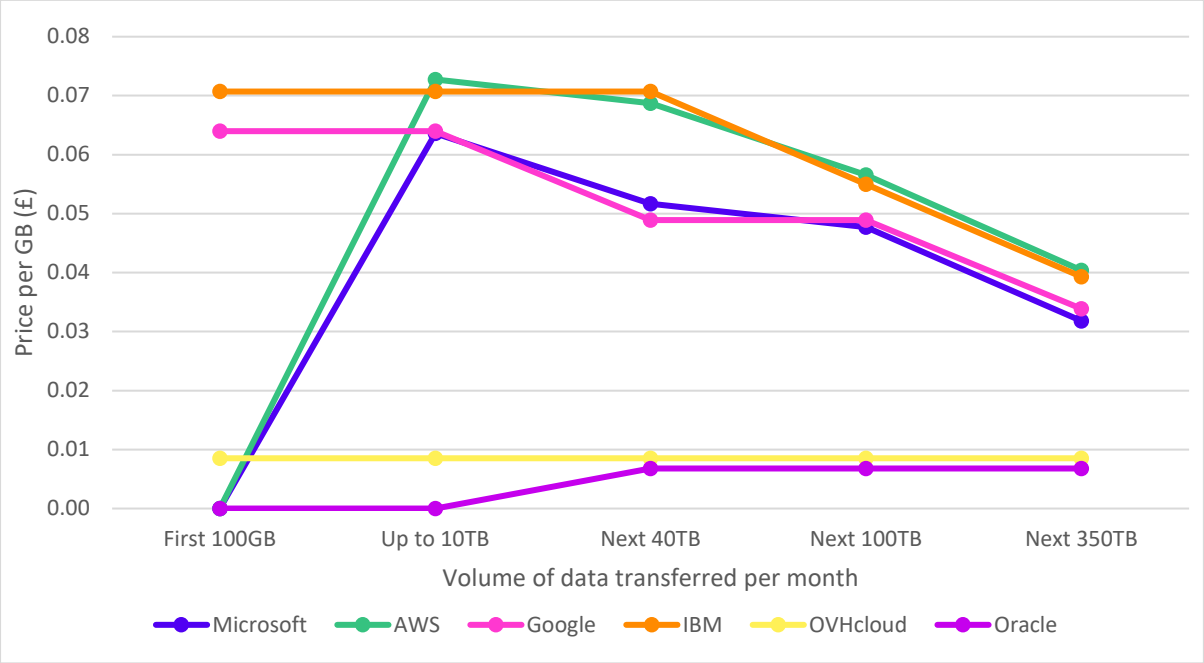


*Sources: Microsoft website. Pricing – Bandwidth | Microsoft Azure; AWS website. Amazon S3 Simple Storage Service Pricing - Amazon Web Services and EC2 On-Demand Instance Pricing – Amazon Web Services; Google website. All networking pricing | Virtual Private Cloud | Google Cloud; IBM website. Cloud Object Storage - IBM Cloud; and Oracle website. Cloud Networking Pricing | Oracle United Kingdom [all accessed 21 September 2023]. OVHcloud response dated 17 November 2022 to s.174 notice of 27 October 2022, question 9, page 13.*

*Notes: Prices shown are for Microsoft: internet egress routed over the internet (ISP network) from North America, Europe to any destination, for Google: Standard Tier egress, and for AWS: Data transfer out from S3 or EC2 to Internet. Prices for GCP converted to per GB from per GiB. Prices for AWS, GCP and OVHcloud converted to GBP based on average exchange rate Jan-Dec 2022. OVHcloud egress charges only apply to data transfer out of object storage. OVHcloud does not charge egress fees for its compute services. IBM structures its egress fees differently to other providers, and so are less directly comparable – IBM egress fees shown in the figure relate to cloud object storage from the Europe - London region.*

5.114 Additionally, we understand that some egress via the public internet will not be subject to these prices:

a) Some customers have negotiated private discounts on egress fees. We discuss this in more detail below.
b) Some cloud providers have agreed to reduce or waive egress fees when customers transfer data between them using Cloudflare's network as part of the "Bandwidth Alliance."[501]

[501] We understand that Google Cloud and Microsoft Azure offer a reduction in egress fees for customers using Cloudflare's network ([✂]% for Google Cloud, [✂] for Azure), and all other providers in the Bandwidth Alliance do not charge egress fees. AWS is not part of the Bandwidth Alliance. For more information on the Bandwidth Alliance, including which cloud providers are included, see Bandwidth Alliance | Reduce Data Transfer Fees | Cloudflare. Also see Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂].

# Egress fees are a barrier to the adoption of multi-cloud

## Data transfer is important for integrated forms of multi-cloud

5.115 Egress fees are likely to be a barrier to the more integrated forms of multi-cloud where a significant amount of data needs to be transferred between the services hosted on different clouds. Transferring data between different services is important to taking services from multiple clouds in an integrated cloud architecture. These scenarios include:

a) moving data between storage, data warehouses and lakes, analytics and machine learning platforms as part of a wider data processing pipeline within an organisation; and
b) exchanging data between applications and their components as part of a service-oriented software ecosystem.

5.116 In addition, most pieces of information within an organisation need to be backed up and archived (e.g. for audit purposes), which requires transferring large volumes of data to long-term storage facilities.

## Customers are concerned by egress fees in relation to multi-cloud

5.117 The market research found that a majority of cloud customers (55%) cite egress fees as a concern. This level of concern is consistent across users of all the major providers, and across customers who are currently using a single cloud provider, and those which are using multiple cloud providers.[502]

5.118 The market research also found that moving data across cloud providers was the biggest single potential challenge to multi-cloud (45% of respondents), for both customers currently using a single cloud provider and those which are using multiple cloud providers.[503]

5.119 In their responses to the interim report, and stakeholder engagement following the publication of the interim report, several customers have mentioned concerns about egress fees affecting their ability to multi-cloud.[504] A number of other respondents to the interim report agreed that egress fees act as a barrier to multi-cloud,[505] including for innovative cloud uses cases and resilience solutions.[506] A couple of stakeholders noted that while egress fees may act as a barrier to multi-cloud, they were not the main barrier.[507]

## Many customers are seeking alternative cloud architectures to minimise or avoid egress fees

5.120 Where services require data to be transferred between clouds, customers will need to pay egress fees. When these services are taken from the same cloud provider instead, customers

---

[502] Context Consulting research data tables, question 63; Context Consulting research report, slide 131.
[503] Context Consulting research data tables, question 31; Context Consulting research report, slide 79. Also see Figure A3.4 and A3.5. A hyperscaler ([✂]) said that the survey research included evidence that some customers conflated egress fees in relation to moving data to other cloud providers with other types of egress fees. While this may be the case for some survey respondents, we do not consider that this undermines this survey finding. See [✂].
[504] [✂] confidential response to interim report [✂], [✂]; BT Group response to interim report, page 2; Gener8 response to interim report, page 2; Ofcom / [✂] meeting, [✂]; Vodafone response to interim report, page 3.
[505] Cloudflare response to the interim report, page 1, [✂] confidential response to interim report, page [✂]; [✂] response to interim report, page [✂], Federation of Communication Services response to interim report, page 3, [✂] response to interim report, page [✂].
[506] BT Group response to interim report, page 2.
[507] See [✂] response to interim report, [✂] and Ofcom / [✂] meeting, [✂].

will not have to pay egress fees.[508] Therefore, egress fees can incentivise customers to design their cloud architectures in such a way as to avoid transferring data between different clouds, and act as a barrier to adopting an integrated multi-cloud architecture. For example, either by purchasing all their services to run on one infrastructure (single cloud), or by siloing services on different clouds (siloed multi-cloud).

5.121 Several customers have explained the impact that egress fees can have on their choices by incentivising the purchase of multiple cloud services through the same provider and limiting the ability to operate a multi-cloud architecture.[509] We have also heard specific examples of customers using a single-cloud architecture, rather than purchasing multiple cloud providers, in part because of egress fees. For example, a customer [✂] stated that avoiding egress charges were a main driver of its decision to purchase multiple cloud services from the same provider.[510] We are also aware of specific examples of customers switching workloads requiring data exchange from a multi-cloud to a single-cloud architecture in part to avoid paying egress fees, even where this means selecting a less preferred solution. These are set out in Box 5.5 below.

**Box 5.5 Examples of customer decisions to move workloads from multi-cloud to single-cloud solutions being influenced by egress fees**

[✂] is moving its data analytics services from Google Analytics to AWS in part to avoid egress costs, even though it considers that AWS's service capability remains below that of Google.[511]

[✂] explained that data egress is the most pertinent obstacle to multi-cloud, considering that data egress has influenced its architectural choice [✂]. It has moved its [✂] to align with its choice of [✂], in part to avoid excessive egress fees.[512]

5.122 Even where customers are choosing to use multiple clouds, many are doing so by siloing workloads on different clouds, rather than adopting more integrated forms of multi-cloud.[513] In some cases, this is motivated in part by a desire to avoid egress fees. For example, [✂]

[508] They may need to pay internal data transfer fees if the transfer involves moving data between geographic regions, but these are typically lower than egress fees.
[509] This includes broadcasters and telecoms operators such as ([✂], BT Group and Vodafone) and other customers such as ([✂]). See BT Group response dated 27 January 2023 to the s.174 notice dated 1 December 2022, question 18, question 20; [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂] question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]; Vodafone response to the CFI, page 3, questions 4.2 and 4.7; [✂] response dated [✂] to our customer questionnaire, question [✂]; and [✂] response dated [✂] to our customer questionnaire, question [✂].
[510] [✂] response dated [✂] to the s.174 notice dated [✂] question [✂].
[511] [✂] response dated [✂] to the s.174 notice dated [✂], page [✂].
[512] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂], pages [✂].
[513] Our market research finds that the largest proportion (45%) of respondents who use more than one IaaS/PaaS used them for different workloads, and that a further 40% spread similar workloads across IaaS/PaaS providers. See Context Consulting market research report, slide 76. We expect that all of the first category, and some of the second category, are undertaking siloed multi-cloud. Further, while some respondents used several cloud providers, very few told us that they were using integrated multi-cloud. We discuss our interpretation of the evidence regarding the frequency of different types of multi-cloud architecture in Annex 3, paragraph A3.17 – A3.29.

noted that individual workloads are located within a single cloud in part to avoid egress fees.[514]

## Egress fees can create uncertainty for customers

5.123 Egress fees can create uncertainty and risk for customers. They can make it difficult to compare the cost of using a single provider compared to the cost of hosting solutions across multiple cloud providers. Even if the egress fees which apply to a particular data transfer are clear and transparent, the overall egress fee payment also depends on the volume of data the customer transfers out. Customers may not be able to perfectly predict how much data they may ultimately need to transfer across different clouds.

5.124 This can affect customers of all sizes. This can be particularly acute for startups, as it is difficult to predict their future business growth which is likely correlated to future data traffic.[515] Smaller customers with less resources may not have the depth of application architecture and cloud infrastructure expertise to understand line-by-line charges.[516] Unpredictability caused by egress fees can also be a concern for larger customers, due to the volume of data flows.[517] The uncertainty of both predicting and tracking the volume of data egress would have a significantly lower financial impact if all services were purchased on the same cloud.[518]

5.125 AWS and Microsoft have noted that they offer tools to assist customers in predicting their costs.[519] Microsoft said that its customers are typically sophisticated and will be able to utilise the tools available to them (including the expertise of third-party consultants) and can assess the cost of ownership across the lifetime of the cloud services (including data egress).[520] However, we note that [✂].[521] We consider customer challenges to predicting cloud spend more generally in paragraphs 5.260 – 5.276 below.

## Some respondents, including hyperscalers, submitted that egress fees are not a barrier to multi-cloud

5.126 Some respondents, including hyperscalers, did not agree with our assessment of the role of egress fees as a barrier to adoption of multi-cloud. AWS told us that egress fees do not restrict customers from accessing multi-cloud solutions as multi-clouding is prevalent, based on its own data, research from Flexera and Gartner, and the findings of Ofcom's market research.[522] AWS also stated that an efficient and well-designed multi-cloud strategy minimises the volume of data transfer out (and thus the associated cost), reducing latency, synchronisation problems, data errors and security vulnerabilities.[523] Microsoft also stated

---

[514] [✂] response dated [✂] to our customer questionnaire, question [✂].
[515] Ofcom / [✂] meeting, [✂].
[516] Cloudflare response to interim report, page 2.
[517] The Register, NASA to launch 247 petabytes of data into AWS – but forgot about eye-watering cloudy egress costs before lift-off, accessed 17 August 2023, and BT Group response to interim report page 2 and Ofcom / BT Group meeting, 5 June 2023.
[518] The customer may still face internal data transfer charges if using a single cloud for all its needs – but these are typically lower than egress fees and are often related to data movement between geographies, rather than between services within a geography.
[519] Including 'AWS Pricing Calculator', 'AWS Cost Explorer', 'Azure Cost Management + Billing' and 'Usage and estimated costs'. See [✂] and [✂].
[520] Microsoft response to interim report, paragraph 235, page 57.
[521] [✂].
[522] AWS response to interim report, paragraph 10, page 5 and [✂].
[523] [✂].

that customers were unlikely to design multi-cloud strategies requiring significant and regular data transfers out for various reasons other than egress fees, such as reducing latency, synchronisation problems, data errors and security vulnerability.[524] A respondent [✂] notes that while egress fees are a factor in ensuring effective competition, and customer choice, focusing on them solely will not resolve the larger 'data gravity' challenge where applications tend to cluster around where their data resides due to bandwidth and latency.[525] Another respondent [✂] argued that multi-cloud was not viable for all but the largest customers due to the resources involved.[526] AWS noted that around 90% of its customers that transferred data out to the Internet pay no egress fees because of its free tier.[527] Microsoft noted that its free tier encourages customers to build and run multi-cloud scenarios involving some data ingress and egress without incurring bandwidth charges.[528] A hyperscaler [✂] said that the evidence did not support a finding that egress fees are the main barrier to multi-cloud adoption.[529]

5.127 We address the evidence on the current prevalence of multi-cloud in Section 4, paragraphs 4.51 – 4.57 and Annex 3. In summary, our analysis indicates that while many customers use more than one IaaS/PaaS provider, the use of integrated multi-cloud is uncommon. Where customers have multi-cloud architectures, this is mainly where a secondary provider is used for siloed use-cases and/or to back-up certain workloads. As discussed in Section 6, our evidence suggests that only some workloads can be siloed and, consistent with this, we see evidence of customers typically concentrating spend around a primary provider and only taking a small number of niche services from a secondary provider.

5.128 While greater integration across clouds may not be needed or desired by every customer, our evidence indicates that there is scope for greater uptake of more integrated multi-cloud solutions by customers. In this regard, we note that customers transfer significant volumes of data internally, implying that they have integrated cloud architectures with data being transferred between different services within a single cloud. For example, the services which currently incur the most data transfers by UK customers (✂) are (✂), and (✂).[530] Moreover, the qualitative part of our market research found that for most customers integrated multi-cloud is the desired model, but the challenge of making multiple clouds work in an integrated way is an obstacle.[531] Similarly, as set out above, our quantitative market research identified that moving data across clouds was the biggest single potential challenge to multi-cloud (45% of respondents), and we have seen examples of customers being deterred from taking services from multiple cloud providers in part due to concerns about the magnitude of potential egress fees.

5.129 In terms of free tiers for egress, observing that many customers do not exceed their free tier allowance does not imply that the free tier allowance is sufficient for them. For very small

[524] Microsoft response to interim report, paragraph 239, page 58.
[525] [✂] confidential response to interim report, question [✂].
[526] [✂] confidential response to interim report, page [✂].
[527] AWS response to interim report, paragraph 17, page 8.
[528] Microsoft response to interim report, paragraph 238, page 57.
[529] [✂].
[530] Source: Ofcom analysis of s.174 notices. [✂] response dated [✂] to s.174 notice dated [✂] question [✂] and [✂]response dated [✂] to s.174 notice dated [✂] questions [✂]; and [✂] response dated [✂] to our s.174 notice dated [✂], Part [✂], question [✂]; [✂] response dated [✂] to our s.174 notice dated [✂], Part [✂] question [✂]; [✂] response dated [✂] to our s.174 notice dated [✂], Part [✂] question [✂]; and [✂] response dated [✂] to our s.174 notice dated [✂], Part [✂] question [✂].
[531] See Annex 3, paragraph A3.28.

customers, this may be the case. But customers remaining within their free tier may be indicative of the importance to customers of avoiding egress fee payments, and how customers are altering their behaviour to ensure they minimise the egress fees which they face. When considering where to place an additional workload, the relevant price for the customers is the marginal price of egress associated with that workload. If customers are already fully utilising their free tier, or fear that additional services spread across different cloud providers would lead them to exceed the free tier for data transfer out, then they are likely to be discouraged from increasing the extent of their multi-clouding where data transfers would be involved.

5.130 Finally, respondents, including the hyperscalers, have pointed to a range of solutions which could mitigate the impact of egress fees on the ability of customers to adopt multi-cloud solutions.[532]

5.131 Hyperscalers pointed to additional services offered to facilitate data transfers out which they said could offer cost savings:

a) [✂] pointed to direct connectivity solutions[533] such as [✂], which do not use the public internet but instead connect to a customer's on-premises network via a dedicated network connection;[534] and
b) [✂] pointed to solutions which rely on physical devices to move extensive volumes of data in and out of the cloud without relying on any network (such as [✂]).[535]

5.132 Further, a hyperscaler [✂] pointed to a solution ([✂]) which it considers minimises the need to continuously transfer data between two providers while still enabling customers to adopt multi-cloud solutions.[536] A hyperscaler ([✂]) pointed to a solution ([✂]) which it considers eradicates the need to perform wholesale migration or large-scale transfers of data between two providers while still enabling customers to adopt multi-cloud solutions.[537] In addition, a hyperscaler ([✂]) anticipates that continued innovation in multi-cloud will further reduce cross-cloud egress.[538]

5.133 We do not consider that these solutions are currently likely to be sufficient to remove the barrier to multi-cloud caused by egress fees for most customers.

5.134 First, direct connectivity solutions from hyperscalers are only likely to be financially viable for large customers, given the additional upfront cost. This is typically determined by the bandwidth of the direct connection and the amount of time for which a port is provisioned (which is charged irrespective of whether data is passing through the port or not).[539] These

---

[532] An individual respondent highlighted "zero-egress cloud object storage services" and managed service providers creating solutions to minimise data transfer costs for their customers, which use a data abstraction layer and caching to store frequently accessed data close to the user's application and reduce roundtrips to the cloud using low-cost direct connections or peered interconnects. Priyank Chandra response to interim report, Q5.1, page 4.
[533] Direct connectivity solutions from hyperscalers are designed to be more scalable and secure while also offering higher bandwidth and better performance.
[534] [✂].
[535] [✂].
[536] [✂].
[537] [✂]. [✂].
[538] [✂].
[539] Dedicated Network Connection - AWS Direct Connect Pricing - Amazon Web Services (accessed 22 August 2023); Pricing - ExpressRoute | Microsoft Azure (accessed 22 August 2023); Pricing|Cloud Interconnect|Google Cloud (accessed 22 August 2023).

costs are likely to be large for most customers.[540] We consider that the cost saving suggested in relation to the use of products such as ([✂])[541] is relatively small, and so is unlikely to be sufficient to address the main issue that these costs inhibit the adoption of integrated multi-cloud solutions.

5.135 Second, innovations which remove the need to egress data between cloud providers (such as [✂]), are currently only available for a limited number of use-cases and will not remove barriers to multi-cloud in relation to using other services. Moreover, the example provided relates to a [✂], therefore suggesting that [✂]. Further, while such innovations may become more prevalent in the future, this is uncertain, especially given the high technical cost of enabling such solutions, and that the hyperscalers (in particular AWS and Microsoft) might have a weaker incentive to introduce them.

## The impact of egress fees as a barrier to adoption of multi-cloud for different customer groups

5.136 The extent to which egress fees act as a barrier to multi-cloud is likely to depend on several aspects of a customer's individual circumstances. These may include how much a customer could save by using a single provider (which will depend on the monthly cost of egress it might face if taking services from multiple providers) and the ability of a rival cloud provider to offset any material egress fees to make it viable for a customer to host integrated cloud services across several clouds. The overall monthly cost of egress it could face will depend on the relevant price it would pay for transferring data from one cloud to another (in particular – the marginal price it would face for transferring more data) and its expectation of the volumes of data it might need to transfer between those services.

5.137 The marginal price of transferring data out of cloud is not uniform. As set out in Figure 5.4, the price of egress varies based on the volume of data being transferred out by a customer in a given month. It also varies between customers based on whether they have been able to negotiate a discount on egress prices, and the magnitude of any discount they have been able to negotiate.

### A small number of customers negotiate discounts on the price of egress, but this is unlikely to be an option for most customers

5.138 We understand that some customers are able to negotiate private discounts for egress fees from the hyperscalers.[542] Evidence received from one hyperscaler suggests that the magnitude of these discounts compared to the list price of egress is large on average for [✂] customers ([✂]%) when excluding [✂], and is even larger on average ([✂]%) when

[540] In addition, in a multi-cloud architecture, a customer may need to purchase direct connectivity from both providers, increasing this cost.
[541] [✂].
[542] [✂], [✂] response dated [✂] to s.174 notice dated [✂], [✂] and [✂] response dated [✂] to our customer questionnaire, question [✂].

including [✂],[543] and can be as high as [✂]%.[544] These discounts can reduce the marginal price of egress to very low levels. Overall, in 2022, customers of the hyperscaler paying list prices (excluding [✂])] paid an average of [✂] per GB when excluding [✂], ([✂] per GB when including [✂]), whereas customers who had negotiated egress discounts paid a much lower average price of [✂] per GB when excluding [✂] ([✂] per GB when including [✂]).[545] For such customers, the marginal price of additional egress will be very low.

5.139 In negotiating an egress discount, customers may take into account their expected egress – including any expected data transfer out as part of a planned multi-cloud architecture. Where they can achieve such a discount, egress fees are much less likely to prevent them from being able to adopt their desired multi-cloud architecture. Therefore, discounted egress fees are likely to be sufficiently low for them not to impact decisions about whether to take services from multiple cloud providers, unless there are restrictions on when those discounted prices apply.[546]

5.140 However, evidence from the hyperscaler suggests that only a very small number of customers pay discounted prices for egress (less than [✂]%) of [✂]).[547] The remaining [✂]% pay the list prices for egress or [✂].[548] We estimate that customers who do not receive egress discounts are competitively important, accounting for over [✂]% of the hyperscaler's cloud revenue.[549]

5.141 That evidence also suggests that, of those very small number of customers that are able to negotiate egress fee discounts, the proportion of larger customers that receive an egress discount is greater than the proportion of smaller customers that receive a discount. [✂]% of the largest [✂] customer [✂][550] pay discounted egress fees, compared to [✂]% of those with a [✂] and compared to [✂]% of the smallest [✂] customers (those with a [✂]).[551] This suggests that most customers (especially smaller customers) may lack the bargaining power to be able to attain a discount on egress fees and may lack a clear understanding of the amount of data they will transfer.

---

[543] [✂]. Ofcom analysis of additional data supplied by a hyperscaler [✂]. [✂]. To calculate this, we consider the average price paid by customers with negotiated discounts for egress in each customer group (where customers are grouped based on their total monthly spend on [✂] services), and we calculate the difference between this price and the average list price paid by customers in the same customer group. This does not adjust for volume effects (the fact that customers receiving negotiated discounts for egress do not transfer the same average data volumes as customers that pay list prices, which matters because the average list price paid will be lower for higher data transfer volumes). However, even if we were to re-calculate the average discount assuming that all customers were paying the minimum possible list price for egress, the estimated discounts would still be large on average ([✂]%).

[544] A large content provider [✂] has negotiated an [✂]% discount on outbound data transfer from [✂]; [✂]% for data transfer out of [✂], [✂]% for [✂]; and a price of [✂] for [✂]. [✂] response dated [✂] to s.174 notice dated [✂], question [✂].

[545] Ofcom analysis of additional data supplied by a hyperscaler [✂]. [✂].

[546] [✂] told us that there are some restrictions on the cloud organisations that are covered by the contract. See [✂] response dated [✂] to s.174 notice dated [✂], question [✂].

[547] Ofcom analysis of additional data provided by a hyperscaler [✂]. [✂].

[548] [✂]

[549] This estimate is based on mapping the proportion of [✂] customers receiving an egress discount in each customer spend band with the total revenue received from UK customers in each of those spend bands. Ofcom analysis of [✂] response dated [✂] to s.174 notice dated [✂] and [✂].

[550] [✂].

[551] Ofcom analysis of additional data supplied by a hyperscaler [✂]. [✂].

5.142 In principle, any customer that plans to transfer large volumes of data out of the cloud on a regular basis would have an incentive to negotiate a discount on egress fees. This will include both customers seeking to deploy multi-cloud architectures with a degree of integration across clouds and customers whose business model relies on transferring data to end users. However, whether a customer has the ability to negotiate a discount on egress fees is likely to depend on their bargaining power, which is a function of how important their overall spend is to the cloud provider, their outside options, and the degree of certainty in the volume of data they expect to egress (which itself requires an element of foresight). Larger customers may be in a better position to negotiate discounts than smaller customers, as they account for the most revenue and may find alternative data transfer methods (such as those in paragraph 5.131 and 5.132) more viable.[552]

5.143 The hyperscalers may also have stronger incentives to negotiate with customers if they are able to identify that customers are likely to be using egress to transfer data to end users, and if that is a key part of the customers' business model. This could include content providers, financial institutions and media businesses. Further, such customers are likely to have greater certainty over their expected data transfer volumes than customers who are looking to use more integrated forms of multi-cloud, and so may be more able to negotiate on pricing and offer commitments related to data transfer. As noted above, egress fees apply both when a customer is transferring data to an end user and when a customer is moving data to a rival cloud provider. Hyperscalers have competitive incentives to make the former easier (and cheaper), and the latter more difficult (and more expensive). However, [✂] told us that they are unable to determine which traffic is for the purposes of switching and which is for other purposes (such as sending data to end users).[553] Therefore, privately negotiated prices for egress may be a way of price discriminating between those who are using data egress to transfer data to end users and those who are using data egress to transfer data to rival cloud providers.[554]

5.144 We do not have enough evidence to fully characterise the customer base receiving discounted egress prices. For the small number of individual customers where we have information, many of the customers which pay low average prices for egress (which may be indicative of receiving discounts) appear to be in industries which may involve high volumes of data transfer to customers (such as broadcasting, entertainment and media and

[552] [✂] states that "customers with high DTO volumes can and tend to enter into private pricing agreements to obtain discounts on the publicly available DTO pricing." [✂].
[553] [✂].
[554] Similarly, we expect that a CDN is most likely to be used by customers looking to deliver content to end users, rather than for transferring data as part of a multi-cloud architecture, especially for customers streaming audio-visual content. Egress fees for transferring data out of CDNs to the public internet, such as AWS CloudFront and Google Cloud CDN, are lower than list prices for egress via the public internet, and, in the case of AWS CloudFront, the amount of data which can be transferred for free each month is ten times larger (1TB per month). See Amazon CloudFront CDN - Plans & Pricing - Try For Free and Cloud CDN: content delivery network|Google Cloud [accessed 30 August 2023].

finance).[555] We are also aware that a large content provider ([✂]) has negotiated a very large discount on egress fees with all hyperscalers.[556]

5.145 We understand that there have been significant reductions in the global per unit cost of AWS providing data egress. It has been claimed that these cost savings have been passed on to customers through lower data transfer out prices.[557] However, to the extent that these cost reductions have been passed on, they appear to have largely been passed on through mechanisms other than list price reductions.[558] Instead, to the extent these cost reductions have been passed on, they appear to have largely been passed on through a combination of the increase in its free tiers (which by definition only applies to a small volume of egress per month) or [✂].

5.146 More generally, there does not appear to be a mechanism whereby those customers negotiating discounts on egress are able to constrain list prices for egress, and so benefit all customers in the market.

## Egress fees from integrated multi-cloud could be material for all customer groups

5.147 The main mechanism through which egress fees could impact competition is by making it more costly to take services from multiple cloud providers. To illustrate the potential for egress fees to disincentivise integrated multi-cloud, we use some stylised scenarios of how much data customers might transfer between providers in a hypothetical multi-cloud architecture. This is set out in Box 5.6 below.

### Box 5.6 Egress fees for hypothetical multi-cloud scenarios

We have considered some illustrative scenarios of the data transfer flows customers might have in a multi-cloud context to demonstrate the egress fees which customers could incur.

How much data customers would transfer in an integrated multi-cloud architecture is uncertain. This will depend on many parameters, such as the type of workloads and architecture the customer desires, whether they choose to duplicate data immediately when created, or transfer data when needed, whether the flow is one-way or two-way, and the growth and maturity of the customer.

---

[555] For example, [✂], all paid less than [✂] per GB of data transferred in 2022, while [✂]. Ofcom analysis of s.174 notices: [✂] response dated [✂] to s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to our s.174 notice dated [✂], question [✂].

[556] [✂] response dated [✂] to s.174 notice dated [✂], question [✂]. We also know that [✂], is one of [✂] largest UK customers. See [✂] response dated [✂] to s.174 notice dated [✂], question [✂].

[557] See Gans, Herve and Masri (2023), Economic analysis of proposed regulations of cloud services in Europe, European Competition Journal, 19(3), Figure 7, page 548-549 and AWS response to interim report, paragraph 15, page 7.

[558] For example, we note that the UK, EU and US list prices for egress above 100GB from AWS Regions to the internet have not changed since 2014, and for data transfers above 50TB have remained unchanged since 1 July 2011, although AWS did increase the free tier from 1GB per month to 100GB per month in December 2021. AWS egress fees for volumes between 100GB and 50TB per month have remained unchanged since 1 December 2014. See AWS website, AWS Free Tier Data Transfer Expansion – 100 GB From Regions and 1 TB From Amazon CloudFront Per Month | AWS News Blog, AWS Data Transfer Price Reduction | AWS News Blog (amazon.com) and AWS Lowers its Pricing Again! – No Inbound Data Transfer Fees and Lower Outbound Data Transfer for All Services including Amazon CloudFront | AWS News Blog. [accessed 29 September 2023].

Our interim report considered three stylised scenarios based on different workloads to illustrate this. Since our interim report, we have received more information on how much data customers have stored in cloud, across six different levels of customer spending. This allows us to illustrate how different types of customer might be affected by egress fees.

We consider that a plausible predictor of the amount of data which a customer would transfer out in a multi-cloud architecture in a given year may be the volume of new data generated by a customer in a year – as one form of multi-cloud architecture would be to replicate any data onto both clouds when it is newly created. We note below (see paragraphs 5.159-5.160) that the overall amount of data stored by UK customers in [✂] clouds grew by around 60% in 2022, and that the overall amount of data stored in the cloud is likely to continue to grow over time. We therefore consider three illustrative scenarios for each customer group, considering scenarios where a customer's data storage grows by 20% ('scenario 1'), 60% ('scenario 2') and 100% ('scenario 3').[559] Scenario 3 also gives a sense of the one-off cost of switching – as it models a customer moving all of their data from one provider to another in one year – which we discuss in more detail in paragraph 5.156 below.[560]

We use list prices to calculate the egress fees which such data transfers would incur on an annual basis, as shown in Table 5.7.[561] We then consider how these compare to an estimate of the annual spend for each of the customer groups in each of these three scenarios.[562] This is shown in Figure 5.8 below.

**Table 5.7 Annual egress fees incurred in illustrative data transfer scenarios by each customer spend band (£)**

| Customer spend band ($) | Scenario 1 | Scenario 2 | Scenario 3 |
|---|---|---|---|
| **Less than 10k** | [✂] | [✂] | [✂] |
| **10k-1m** | [✂] | [✂] | [✂] |
| **1m-5m** | [✂] | [✂] | [✂] |
| **5m-10m** | [✂] | [✂] | [✂] |
| **10m-20m** | [✂] | [✂] | [✂] |
| **More than $20m** | [✂] | [✂] | [✂] |

*Notes: Data related to customer spend bands were provided in USD. However, egress fee figures reported in the table have been converted to GBP. Figures have been rounded to the nearest thousand, unless the figure is less than one thousand.*

[559] We calculate the data transfer volumes for these illustrative scenarios based on [✂]. In these illustrative scenarios, we assume that a customer transfers data out in equal amounts in each month. This matters because egress fees, and associated free tiers, are based on the volume of data a customer transfers out each month.

[560] When switching, a customer may move its data across in a shorter time period than one year (although switching may require several months of integrated multi-cloud). Moving the same volume of data across in a shorter time period would lead to a different total egress fee payment.

[561] Using [✂].

[562] Using [✂].

**Figure 5.8 Egress fees in illustrative data transfer scenarios as a percentage of annual spend for each customer spend band**

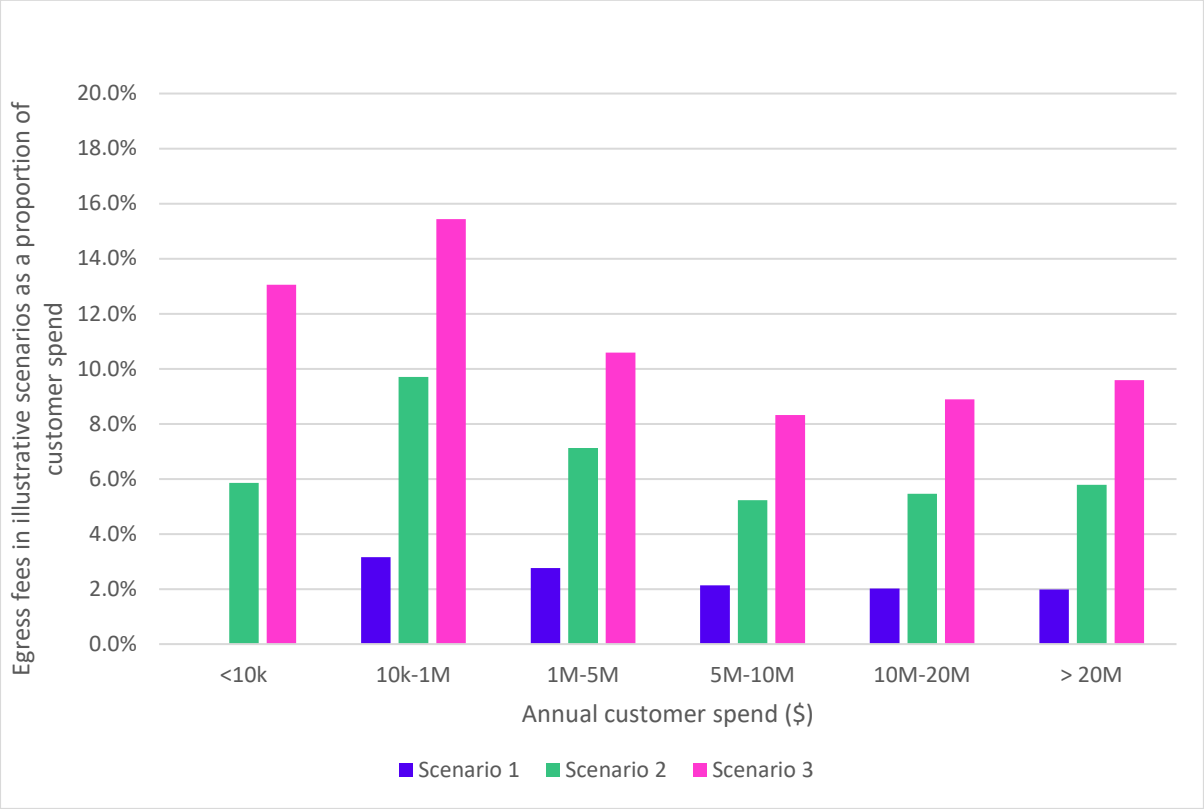


5.148 These scenarios indicate that egress fees could significantly increase the cost to customers of taking a service from a rival cloud provider. These costs could range from a few thousand pounds for small customers, to over £1m for the largest customers. Figure 5.8 shows that egress fees in these scenarios could be a material fraction of total cloud spend for all customer groups.[563] Customers could face egress fees of between 5-10% of annual customer spend in Scenario 2, and up to 15% of annual customer spend in Scenario 3.

5.149 Further, egress fees may significantly increase the cost of taking a particular service from a different cloud provider, and in some cases could double the effective price paid by a customer. For example, it would cost a customer around £400 per month to analyse a 9TB data set stored in Google BigQuery.[564] By comparison, a customer with data stored in AWS or Microsoft clouds would need to pay around £600 per month to transfer 9TB of data to Google Cloud (at list prices), in addition to the cost of analysing the data in BigQuery.[565] To analyse a 210TB data set stored in Google's cloud in BigQuery would cost around £9,000 per month, whereas egressing 210TB of data in a month from either AWS or Microsoft at list prices would cost an additional c.£10,000 per month.[566] This suggests that egress costs may

[563] Very small customers may find that the free tier is sufficient, even if they were to adopt an integrated multi-cloud architecture.
[564] The PAYG price of Google BigQuery consists of a storage price (which is no more than £0.04 per GB) and a compute price (which is c.£0.006 per GB). See Pricing|BigQuery: Cloud Data Warehouse|Google Cloud [accessed 21 September 2023], converted into GBP per GB.
[565] In this example, we abstract from any internal data transfer charges which a customer may face when storing and analysing its data in the same cloud, but, in any case, these are typically lower than egress fees.
[566] The volumes of data transfer used in these examples correspond to Scenario [✂] for customers in the $[✂] and $[✂] customer spend categories.

be hard for rival cloud providers to compensate for through a discount, particularly if a potential secondary provider is competing for only a few workloads.

5.150 These are purely intended to be illustrative scenarios. However, other indicators suggest that they could provide a plausible estimate of the potential egress costs associated with multi-cloud architectures. We note that Cloudflare has found that reducing or eliminating egress fees has saved customers using the Bandwidth Alliance between 7.5% and 27% off their monthly bill, the lower bound of which is similar in magnitude to our estimates of the proportion of annual spend in Scenarios 2 and 3.[567] We also note that, for the four largest customer groups, the data transfer volumes assumed in Scenario 2 are broadly equivalent to [✂].[568]

5.151 Our scenarios focus on customers paying the list prices for egress. For customers which have negotiated discounted prices for egress, egress fees in these scenarios would be much less material, both in absolute and relative terms. But as we set out above, only a very small number of customers pay discounted prices for egress.

## Egress fees can act as a barrier to switching

5.152 Many respondents to our interim report said that egress fees can act as barrier to switching.[569]

5.153 Nearly a quarter of respondents to the market research said data charges are a challenge to switching – although only 6% said it is the biggest switching challenge.[570] The proportion of customers citing egress fees as a concern in our market research is higher for those who have switched, or considered switching provider, than for those who have not.[571] We have also heard that egress fees when switching could be particularly large in relation to the transfer of data warehouses.[572]

---

[567] See Cloudflare response to interim report, page 1 and Empowering customers with the Bandwidth Alliance (cloudflare.com) [accessed 7 September 2023]. Cloudflare estimated the potential savings for a sample set of its customers using a $0.08 per GB retail price compared to a discounted $0.04 per GB for large amounts of data transferred, compared to the amount of money those customers spend on Cloudflare. It also noted that customers could save more by using one of its partners with whom the cost is $0 per GB.

[568] [✂].

[569] BT Group response to interim report, page 2; Cloudflare response to interim report, page 3; Federation of Communication Services response to interim report, dated 17 May 2023, page 2 and page 4; Gener8 response to interim report, pages 2 and 3, [✂] confidential response to interim report, question [✂]; [✂] confidential response to interim report, pages [✂]; [✂] confidential response to interim report, page [✂] and Ofcom / [✂] meeting, [✂]; Vodafone response to interim report, page 3; [✂] response to interim report, question [✂], page [✂].

[570] This is the sixth largest switching barrier. Interoperability appears to be a greater concern for customers seeking to switch cloud providers. However, the most mentioned main challenge of switching (20% of respondents) was the time and cost of making the change – which may include data charges – and so this 6% figure may understate the proportion of customers which considered data charges to be the main challenge. Context Consulting research data tables, question 52 and 53. We also note a survey by Public First, which found that 15% of "switchers" said that the majority or vast majority of costs when switching were external (such as data transfer fees), and that 15% of IaaS/PaaS users said that the majority or vast majority of costs when switching would be external. See questions 34 and 37 of the Public First's survey (available at: https://www.publicfirst.co.uk/files/CCIA_Survey.xlsx [accessed 7 September 2023]). Also see paragraph A3.81 and A3.82.

[571] Context Consulting research data tables, question 63; Context Consulting research report, slide 131.

[572] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

5.154 Other respondents to our interim report, including the hyperscalers, said that the one-off cost of egress associated with moving a customers' data from one cloud provider to another is relatively small compared to customer's overall spend on cloud services, and that such a low one-time cost is unlikely to deter customers from switching.[573] Google also said that egress fees only account for a small proportion of an average customer's total spend [✂] and the proportion of the total relating to switching/exit is therefore even smaller.[574] Some respondents more generally considered that egress fees were not a material barrier to switching, especially in relation to other barriers such as technical barriers and the complexity of switching.[575] Microsoft also referred to the results of our survey as indicating that egress fees are not a significant barrier to switching because only a small proportion of respondents said that egress fees constituted the largest challenge to switching between cloud infrastructure providers.[576] [✂] has also referred to the results of a survey by Public First which it claimed showed that customers did not see data transfer-out fees as inhibitive to switching.[577]

5.155 The 'one-off' egress cost associated with transferring all of a cloud customer's data from one cloud provider to another is likely to be a relatively small proportion of the overall cost of the switch. For some customers this cost could be a relatively significant portion of their current annual spend – and potentially millions of pounds for the very largest customers.[578] However, unlike when adopting an integrated multi-cloud architecture, this cost would only be incurred once, rather on a repeated basis for a prolonged period.

5.156 We now understand that the one-off egress cost is likely to understate the total egress cost associated with switching. This is because switching both incurs one-off costs of data transfer to the gaining cloud provider and often necessitates a period of integrated multi-cloud running during the switching process. As set out in Section 5, paragraph 5.25, a customer seeking to switch all of its cloud services from one provider to another is likely to do so in a number of stages – rather than switching all its services at once. During the switch, to maintain continuity of service, the customer will have services hosted on multiple cloud infrastructures and will require data to be transferred between them. The customer would be incurring egress fees during this period, which could be many months.

5.157 The impact of egress fees for a switching customer is therefore not just the additional cost caused by the one-off transfer of data from one cloud to another, but also the additional cost (and uncertainty) of egress fees during the period of integrated multi-cloud

---

[573] AWS response to interim report, paragraph 17, page 8; Google response to interim report paragraph 10, page 2 and [✂]; Microsoft response to interim report, paragraph 237, page 60; [✂] confidential response to interim report, pages 4, 6-7; and Priyank Chandra response to interim report, question 5.1, page 4.
[574] Google response to interim report paragraph 10, page 2 and [✂].
[575] [✂] response to interim report, page [✂]; [✂] response to interim report, page [✂]. Also, [✂] noted customers were more focused on technical switching barriers than egress fees. [✂] response to interim report, page [✂].
[576] Microsoft response to interim report, paragraph 236, page 57.
[577] [✂].
[578] Our illustrative Scenario 3 of Box 5.6 provides an estimate of the potential cost to a customer of switching all of their data out – suggesting that the one-off costs could be between 8-15% of a customer's annual cloud spend. Other customers have told us that at present, the financial cost of egress fees were they switch cloud providers could be at least £100,000 – although this may be a relatively small proportion of their current cloud spending ([✂]%). See [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

necessitated by the switch. This increases the egress-related cost of switching, and could be significant for some customers. However, this will remain a temporary effect. Further, it may be more feasible for a customer to secure offsetting discounts on their service prices from the gaining provider when switching, than when adopting a multi-cloud architecture. In particular, where the cost may be a relatively small proportion of the total revenue the new provider expects to receive over the lifetime of the switching customer's contract.

5.158 Overall egress fees may act as a barrier to customers who would have to move a large volume of data, and those that would need to run an integrated multi-cloud architecture for the duration of their switch. But we also acknowledge that our market research found that the proportion of respondents citing "data charges" as a barrier to switching was relatively low. Therefore, while egress fees may act as a barrier to switching, we expect this to affect a smaller fraction of customers compared to egress as a barrier to multi-cloud.

## The impact of egress fees as a barrier to both multi-cloud and as a barrier to switching may grow over time

5.159 The extent of the barrier to both multi-cloud and switching from egress fees reflects the amount of data a customer has stored in the cloud, and the amount of data it seeks to transfer between different services as part of its cloud usage.

5.160 We understand that the volume of data stored in the cloud is growing. In 2022, the amount of data stored by UK customers in [✂] clouds grew by around 60%.[579] This is likely to continue to grow over time, as more customers migrate to cloud, as existing customers migrate more workloads to cloud, and as more data is generated by cloud services.[580] As the diversity and sophistication of a company's IT stack grows, customers may need to move more data between more of a customer's applications.[581] This could lead to egress fees increasing over time if a customer operates a multi-cloud architecture that requires data to be exchanged between clouds.[582] This means that costs associated with transferring data are likely to become an increasingly important part of customer decisions as to which cloud services to purchase, and the extent to which these services are taken from single or multiple providers.

## Hyperscalers' egress fees at their current levels are unlikely to be necessary for cost recovery

5.161 Our interim report considered three indicators which suggested that hyperscalers were setting egress fees above the incremental cost of providing data transfer:

---

[579] Ofcom analysis of s.174 notice responses. [✂].

[580] For example, the market research found that 43% of current users of IaaS/PaaS (or those considering using IaaS/PaaS) reported the intention to migrate more workloads to the cloud in the next 18 months. Context Consulting research report, slide 43; Section 3, paragraph 3.35 and Figure 3.5. Also, IDC, 2021, Future Proofing Storage, Modernising Infrastructure for Data Growth Across Hybrid, Edge and Cloud Ecosystems, future-proofing-storage-wp.pdf (seagate.com), page 8 [accessed 13 September 2023], which finds that the installed base in public cloud is expected to grow by 35.9% per year. Gener8 highlights that storage needs are likely to increase exponentially in the years ahead. Gener8 response to the interim report, question 5.1.

[581] Cloudflare response to the interim report, page 1.

[582] [✂] told us that it expects to see its egress costs rise over time. Ofcom / [✂] meeting, [✂]. Cloudflare said that egress fees almost always increase over time. Cloudflare response to the interim report dated, page 1.

a) the difference in the prices charged by the hyperscalers and other providers such as OVHcloud and Oracle;
b) that prices charged by the hyperscalers are much higher than the likely transit charges they face, which we considered to be the key cost unique to data transfers out; and
c) other aspects of the structure of egress fees charged by the hyperscalers, including use of free tiers.

5.162 Our view was supported by some smaller cloud providers in their response to our interim statement. For example, Oracle noted that AWS's egress fees were far higher than its own egress fees, which it said were set equal to its cost of data transfer.[583] It further stated that the main costs of providing data transfer are generally all transit costs.[584] A cloud provider [✂] stated that egress fees are not technically justified when it comes to transferring data to another cloud service provider.[585] Cloudflare said that the marginal costs of data transfer for the hyperscalers are often near-zero for large customers. It also said that while it recognised fixed costs associated with build-out of networks need to be recouped, these should not be recovered through egress fees.[586]

5.163 However, the hyperscalers stated that the costs of providing data transfer include a broader set of costs than transit, including those relating to the network infrastructure over which data is transferred.[587] AWS said that these costs are significant because of the enhanced security and reliability features in its proprietary network.[588] [✂] said that [✂].[589] As such, [✂].[590] [✂] said that its network presence was far larger than that of a smaller provider ([✂]).[591]

5.164 Egress fees charged by other providers may not be directly comparable with those charged by the hyperscalers, as the scope of the service provided may be different. We also acknowledge that the hyperscalers have invested, and continue to invest, in providing and maintaining high quality, secure networks and that they use their networks to transfer data out of their cloud. As such, it is likely that some of the costs associated with the investment in that network infrastructure might be attributed to the cost of providing egress.

5.165 However, it does not follow that the hyperscalers need to maintain the current level of egress fees to recover the cost of providing egress.

5.166 Google and Microsoft charge higher egress fees than those shown in Figure 5.4 where a customer opts to route data over its private network.[592] As such, the difference between

---

[583] Oracle response to the MIR consultation, page 3.
[584] More specifically, internet provider transit, costs for circuits to connect one DC to another DC, or costs to connect metros. Oracle email to Ofcom dated 9 August 2023.
[585] [✂] confidential response to the interim report, page [✂].
[586] Cloudflare response to the interim report dated, page 2.
[587] AWS response to interim report, paragraph 14, pages 6 and 7; Google response to the interim report, paragraph 9, page 2; and Microsoft response to the interim report, paragraph 229, page 56.
[588] AWS response to interim report, paragraph 14, pages 6 and 7 and [✂].
[589] [✂]
[590] [✂].
[591] [✂].
[592] Both Google and Microsoft offer two egress tiers – one where data is transferred through their own network (Google Premium Tier and Routing via Microsoft global network), and one where data is routed via transit (ISP) networks (Google Standard Tier and Microsoft Routing over public internet (ISP network)). See Routing preference in Azure - Azure Virtual Network | Microsoft Learn and Network Service Tiers overview|Google Cloud. We understand that the majority of AWS egress is routed through AWS's own network. [✂].

their egress fees shown in Figure 5.4 and the fees charged by smaller cloud providers like OVH Cloud and Oracle cannot be fully attributed to their investment in underlying network infrastructure.

5.167 More generally, a provider would expect revenues from specific services to at least cover the incremental costs incurred to provide those services. It would also expect to earn sufficient aggregate revenues across all its services to cover all its costs (including common costs and the costs of any capital invested). In this context, we understand that aspects of the network used for cloud egress are shared with the hyperscalers other significant digital businesses – meaning that there may be scope for common costs to be recovered to a wide range of other services.[593] In practice, a provider will likely exercise some flexibility in how it sets the price for each service, with market conditions more likely to influence pricing decisions than the requirement for the price of each service to cover an arbitrary allocation of the common costs.

5.168 A hyperscaler ([✂]) provided us with an estimate of egress costs for the period [✂], which included network and data centre costs associated with providing egress (and other data transfer services), plus a revenue-based attribution of costs which are common across many services. Even after including this attribution of common costs, this data indicates that [✂], it ([✂]) was earning a positive margin ([✂]%) on egress during this period, indicating that [✂] net egress prices were more than sufficient to recover all the incremental costs of providing egress, and contribute to the shared costs of providing [✂] services.[594] In addition, we note that the common costs are around [✂] of the total allocated egress cost, and they have been allocated across services (including across different data transfer services) based on relative revenue.[595] This suggests the allocation of common costs to egress might be relatively high compared to other data transfer services, because prices for those other services are typically lower than egress. We therefore maintain our view that hyperscalers are currently setting egress fees (i.e. egress list prices) that are likely to be higher than the incremental costs of providing the service.

5.169 This is consistent with our observations on the pricing of data transfer services where market conditions may incentivise providers to set prices closer to incremental costs. For example, hyperscalers do not charge for data ingress (i.e. the effective price is zero) and, as set out above, some customers which transfer a high volume of data out have negotiated large discounts on egress fees.[596]

## Conclusion on egress fees

5.170 We consider that egress fees are likely to act as a barrier to more integrated multi-cloud architectures for most customers. This is due to the significant additional cost and uncertainty they create for customers if they were to migrate data between two providers on a regular basis. Egress fees can lead to a significant saving for hosting all services associated with a use case on a single cloud, and can be significant for customers across different spend levels. This has the potential to impact the choices of a material fraction of

[593] For example, see Google website. The economic advantages of Google Cloud Networking | Google Cloud Blog [accessed 19 September 2023].
[594] [✂]; Ofcom analysis of [✂].
[595] ([✂]) are allocated on the basis of revenue proportion. [✂] has stated that these costs are not directly associated with data transfer. Ofcom analysis of [✂].
[596] [✂].

customers. We acknowledge that egress fees are unlikely to be a barrier to muti-cloud for customers that have negotiated private discounts on egress. However, we estimate that such discounts are limited to a small portion of the market in revenue and customer terms.

5.171 Egress fees may also act as a barrier to switching for customers who would have to transfer large stores of data, and those that would need to run an integrated multi-cloud architecture for the duration of their switch. For other customers, the barrier to switching from egress fees is likely to be weaker than as a barrier to multi-cloud.

5.172 We also find evidence indicating that egress fees at their current level are unlikely to be necessary for cost recovery, and that egress list prices are likely to be higher than the incremental costs of providing the service.

5.173 In making these findings, we also note that other international regulators (such as the *Autorité de la concurrence*, Authority for Consumers and Markets (ACM), and the European Commission) have similarly found concerns in relation to the scope for egress fees to create barriers to multi-cloud and switching.[597]

# Committed spend discounts

## Introduction

5.174 Large customers often negotiate with hyperscalers for committed spend discounts, where the customer commits to spend a minimum amount across the hyperscaler's cloud services over a period of years, and in return, receives an individually negotiated percentage discount on list prices. AWS's Enterprise Discount Program[598] and Microsoft's Azure Consumption Commitment[599] are examples of the discounts which this section discusses.

5.175 In our interim report, we raised concerns that the structure of these agreements and the way hyperscalers use them could create an incentive for customers to buy all or most of their cloud infrastructure needs from a single provider, creating a barrier to multi-cloud. This barrier does not relate to the practice of discounting or of customers commitments in general, which may be desirable features of a well-functioning market. The scope for committed spend discounts to create a barrier to multi-cloud is instead prompted by their structure and specific characteristics. Accordingly, we provisionally concluded that competition for the incremental workloads of customers with existing commitments may be dampened.

5.176 We have received responses to our interim report and have gathered further evidence which has developed our understanding of these discounts and their potential effects in the cloud infrastructure services market. In particular, we have improved our understanding of the degree to which discounts are individualised to specific customers, and of the hyperscalers' practices where they have bargaining power over customers.

---

[597] *Autorité de la concurrence*, 2023. Press release. [accessed 25 August 2023], ACM, 2022. Market Study Cloud services, page 5 [accessed 25 August 2023] and European Commission. Data Act: Commission proposes measures for a fair and innovative data economy [accessed 11 September 2023].
[598] AWS website. Enterprise Customers [accessed 14 September 2023].
[599] Microsoft website. Track your Microsoft Azure Consumption Commitment (MACC) [accessed 14 September 2023].

## The nature of committed spend discounts

5.177 Committed spend discounts relating to specific products or families of products are offered by many cloud providers, including the hyperscalers and non-hyperscalers. They are generally available to all customers, including those who agree to the standard terms of service on the cloud providers' websites.[600] In addition to these, all three hyperscalers offer cross-service committed spend discounts, where spending across all (or almost all) of their cloud products will draw down the commitment, but only as part of privately negotiated agreements with large customers.[601] Our assessment has focused on these cross-service, individually negotiated committed spend discounts and their use by hyperscalers.

5.178 In our interim report, we set out our understanding of how these discounts function, gave some descriptive statistics on their characteristics, and how they are negotiated.[602] In this section, we present our current understanding of these factors.

### Mechanisms of committed spend discounts

5.179 This subsection gives our understanding of how these agreements generally function based on the information we have gathered, although we acknowledge that some terms may be tailored for individual customers in some cases.[603]

5.180 We understand that the discount a customer receives from a spending commitment is set as a simple percentage, and is cross-service, meaning it applies to almost all of the cloud services offered by the cloud provider and that spending on all of these products is eligible to draw down the customer's commitment. Services which are excluded from the spending commitment (meaning that they do not benefit from the discount and that they do not draw down the commitment) are often people-based.[604] The percentage discount will apply to all eligible spending, including that above the customer's commitment.[605] Spending on hyperscalers' marketplaces (including on listed third-party services) is generally eligible to draw down customers commitments, within some limits.[606]

5.181 Such commitments function as minimum spends, where a customer that does not spend up to their commitment for a given period is obliged to pay the difference between their actual

---

[600] See, for example, Google's documentation of its resource-based committed use discounts. Google website. Resource-based committed use discounts [accessed 7 February 2023].
[601] [✂] response dated [✂] to the s.174 notice dated [✂], [✂] question [✂]; [✂] response to the s.174 notice dated [✂], Part [✂] question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], [✂] question [✂].
[602] Ofcom, 2023. Interim report, paragraph 5.127 to 5.134.
[603] Some cloud providers have also changed their policies and practices over time. An example of a customer's committed spend discount agreement with AWS is available online. SEC website. AWS Enterprise Discount Program Addendum [accessed 13 September 2023].
[604] For example, see AWS's list of "ineligible services". AWS website. Ineligible Services [accessed 13 September 2023].
[605] Ofcom, 2023. Interim report, paragraph 5.128.
[606] Two hyperscalers told us that they typically cap the amount of a customer's commitment which may be drawn down through spending on their marketplaces. One of the hyperscalers, which now uses a cap, told us that they previously allowed marketplace spending to draw down commitments at [✂] the rate of spending on their first-party products. Another hyperscaler told us marketplace spending on third-party services draws down customers' commitments if they are substantially platformed on the hyperscaler's own infrastructure. [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂].

and committed spend as a lump sum.[607] Some cloud providers stated that they would attempt to address any issues with customers which are not on track to meet their commitments in advance, for example by negotiating an extension of the customer's contract.[608]

## Characteristics of committed spend discounts

5.182 Many of the discounts offered by hyperscalers are widely available to customers, including some that involve customers pre-paying or committing to spend on specific cloud services (e.g. Reserved Instances and Savings Plans). The individually negotiated cross-service committed spend discounts, which are the focus of this section, are generally only available to larger customers – i.e. those who have some degree of experience in using cloud services,[609] and those spending more than $[✂] per year with a provider.[610]

5.183 Data submitted by two hyperscalers indicates that past this $[✂] per year threshold customers become more likely to have negotiated a committed spend discount the more they spend with them. For customers with more than $[✂] of spend a year, negotiated committed spend discounts are widespread, and they cover the vast majority of their customers with an annual spend above $[✂][611] As such, committed spend discounts affect a large proportion of the cloud infrastructure services market. Hyperscalers have told us that customers with privately negotiated committed spend discounts accounted for [✂]% of their UK cloud customers by count, but for [✂]% [(✂)] a high proportion of their UK cloud revenues.[612]

5.184 The discounts that customers receive across their cloud spending with a hyperscaler in exchange for a commitment can vary widely, including for customers making commitments of similar sizes for similar durations, though they are most commonly between [✂]%.[613] For a small proportion of customers, discounts can exceed [✂]%.[614] One hyperscaler submitted that its average cross-service discount provided to customers in exchange for a given spending commitment had [✂].[615]

---

[607] One hyperscaler told us that contractual damages may apply in some cases. [✂] response dated [✂] to the s.174 notice dated [✂], [✂], question [✂]; [✂] response to the s.174 notice dated [✂], [✂] question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], [✂] question [✂].
[608] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], [✂] question [✂].
[609] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂] and email from [✂], dated [✂].
[610] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]; [✂] confidential response to the interim report, paragraph [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], [✂] question [✂].
[611] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[612] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂].
[613] [✂]. [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂].
[614] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂].
[615] [✂].

5.185 We have seen examples of committed spend agreements with durations between one and seven years.[616] The hyperscalers told us that in 2022 [✂]% of their total UK cloud services revenues were accounted for by customers with spending commitments and agreements lasting at least one year, [✂]% of revenues were accounted for by customers with spending commitments and agreements lasting at least 3 years, and [✂]% of revenues were accounted for by customers with spending commitments and agreements lasting at least 5 years. The hyperscalers vary in their use of agreements longer than 5 years, from [✂]% of revenues to [✂]% of revenues.[617]

## Negotiation of discounts

5.186 Committed spend discounts represent the primary way that large customers are able to negotiate with hyperscalers on price. One hyperscaler submitted figures on the relative sizes of different discount programmes they offer customers, in terms of the percentage impact they have on its total cloud revenue compared to if it sold the same quantity of services at list prices. They show that committed spend discounts are approximately [✂] three times larger than Reserved Instances and Savings Plans, and approximately [✂] 10-15 times larger than cloud credits.[618] On this basis, their submission states that committed spend discounts are "by far the most important category and their share has been significantly increasing."[619]

5.187 A spending commitment may be renegotiated during its contract duration, by the mutual agreement of the customer and cloud provider. We understand that the most common reason for doing so is if a customer expects to exceed their existing commitment and wishes to make a higher commitment in exchange for a higher discount.[620] Otherwise, we understand that customers and hyperscalers will generally renegotiate in the final months of the customer's existing agreement, for an agreement which commences following the end of an existing contract's duration.

5.188 Other contractual terms are often negotiated alongside committed spend discounts. This includes other forms of discount, such as cloud credits and migration-based discounts. One hyperscaler told us that typically only its largest customers, roughly those spending above $[✂] per year, would negotiate single-service discounts which sit alongside their cross-service discount.[621]

## Individualisation of discounts

5.189 An individual customer who is negotiating an agreement that includes a committed spend discount will normally be able to obtain a higher percentage discount by increasing the size of their commitment and the duration of their commitment.[622] However, we noted in our interim report that committed spend agreements are privately and individually negotiated,

---

[616] An example of a seven-year commitment to a hyperscaler was included in [✂] response dated [✂] to our customer questionnaire.
[617] One hyperscaler was only able to provide this information for 2021 rather than 2022. [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂].
[618] [✂].
[619] [✂].
[620] [✂].
[621] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[622] [✂]; and Microsoft response to the interim report, paragraph 158.

meaning there is scope for the discounts customers receive in return for a given commitment to vary.[623]

5.190 One hyperscaler submitted that committed spend discounts are "heavily negotiated" and that [✂]% of their committed spend discount deals provided to customers are not accounted for by [✂].[624] We have also heard a stakeholder account of customers receiving very different discounts from the same hyperscaler for commitments of a similar size, duration and terms.[625]

5.191 Two hyperscalers [✂] submitted data on the committed spend discounts they have agreed with their customers.[626] We have analysed the distribution of discounts which their customers are receiving, compared to the size and durations of their commitments.[627] This analysis is described in greater detail in Annex 5. This data shows that for each of these hyperscalers, there is a wide distribution of percentage discounts being given for similar levels of commitment. For example, one hyperscaler's, ([✂]), data suggests that customers with a commitment spend of [✂] million and a contract duration of around [✂] months can receive discounts ranging from around [✂]%, while [✂] data suggests that customers with a commitment spend of [✂] million and a contract duration of around [✂] months can receive discounts ranging from around [✂]%.[628] Further, the [✂] R-squared statistics in the regression analysis which hyperscalers have submitted shows that, even when controlling for a range of parameters related to customers' commitments, there remains a lot of unexplained variation in the levels of discount received by customers.[629]

5.192 The evidence we have seen indicates that the individualisation of discounts by these hyperscalers is widespread.

## How committed spend discounts may affect customers' decisions

5.193 In our interim report, we explained that the structure of committed spend discounts has the potential to create an incentive for customers to concentrate all or most of their cloud use with a primary cloud provider.[630] We also identified a further concern that competition for the incremental workloads of customers with existing committed spend discounts may be affected.[631] We have gathered further evidence from stakeholders, including customers, hyperscalers, and smaller cloud providers on the way committed spend discounts may affect customer choices, and address each of these in turn.

---

[623] Ofcom, 2023. Interim report, paragraph 5.132.
[624] This is how the hyperscaler terms their baseline schedule of commitments and corresponding discounts which they state they use. [✂].
[625] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[626] In the case of one of these hyperscalers, [✂], the data submitted included various limitations which may make analysis based on it less robust. These limitations are discussed in Annex 5.
[627] As discussed earlier, there other parameters of contracts which are negotiated alongside committed spend discounts, including referenceability and various other forms of discount.
[628] See Annex 5, paragraph 5.31.
[629] See Annex 5, paragraph 5.26.
[630] Ofcom, 2023. Interim report, paragraph 5.162.
[631] Ofcom, 2023. Interim report, paragraph 5.163.

## Incentive to concentrate spend with a primary cloud provider

5.194 A customer's discount across its entire spend will increase as they increase their monetary commitment to a provider during a negotiation or renegotiation. A customer will therefore maximise the size of the percentage discount they receive by placing as much of their cloud spending as possible with a single provider. One hyperscaler told us that its customers typically want their spending commitments to be as close as possible to their anticipated spend with them to earn the maximum available discount.[632] This demonstrates how customers have an incentive to maximise the size of the commitment they make to a single provider to maximise their discount.

5.195 Widespread individualisation of agreements and discounts may further contribute to a customer concentrating their spending with a primary provider. By tailoring the schedule of discounts to customers' product needs and total spending, the hyperscalers may be able to structure discounts to encourage single sourcing.[633] For both of the hyperscalers which submitted data on the discounts they negotiated with customers, we can observe a relationship between commitment size and discount, but we also observe wide variation in the discounts received by committed spend discount customers making similar commitments.[634]

5.196 In our interim report, we provided evidence showing that the decisions of some customers were being affected by their spending commitments. This included accounts from customers that their committed spend discounts had induced them to purchase incremental cloud needs from within the ecosystem of the same cloud provider they were already using. Other customers told us that their commitments to hyperscalers do not represent a barrier to them choosing the most appropriate providers for their cloud needs.[635] This evidence is consistent with committed spend discounts affecting competition for individual workloads, though the extent of this may depend on individual customers' circumstances.[636] In this subsection we provider further accounts on how discounts impact on the purchasing decisions of customers.

5.197 Microsoft submitted that customers' commitments are generally planned by customers based on specific workloads within Azure only, and so do not resemble discounts which would generally raise competition concerns.[637] In our view, the fact that committed spend discounts are cross-service in nature and apply to a customers' entire spend means that they do not simply apply to individual workloads. Moreover, the individualisation of discounts and commitments means that providers may have the ability to structure these discounts to encourage customer to concentrate their spend with a single cloud provider.

5.198 Overall, we are concerned that committed spend discounts create an incentive for customers to source all or most of their needs from their primary cloud provider, which could pose a barrier to multi-cloud. This includes when customers first move into the cloud as well as when they renegotiate their agreements with their primary cloud provider. In the remainder of this section, we first consider the submissions hyperscalers have made on this

---

[632] The hyperscaler submitted that it is not in their interest to have customers making spending commitments that they cannot meet. [✂].
[633] Ofcom, 2023. Interim report, paragraph 5.136(c).
[634] See Annex 5, paragraphs A5.31 to A5.33.
[635] Ofcom, 2023. Interim report, paragraph 5.140.
[636] Ofcom, 2023. Interim report, paragraph 5.138.
[637] Microsoft response to the interim report, paragraph 163.

topic, then we discuss the evidence we have gathered from our engagement with customers and other stakeholders.

*Hyperscalers' evidence of multi-cloud does not address our concern that customers are encouraged to concentrate all or most of their spend with their primary provider*

5.199 Hyperscalers submitted that many of their customers with spending commitments are observed to purchase from multiple cloud providers, and that this shows committed spend discounts do not represent an incentive to single source.[638]

5.200 One hyperscaler, [✂], submitted analysis of its internal data on opportunities won and lost among its customers which have spending commitments, submitting that this shows that [✂]% of the customers represented in this dataset awarded an opportunity to another cloud provider before or during the period they had a commitment to [✂]. This analysis is discussed in more detail in Annex 3.[639] Another cloud provider similarly submitted that many of their customers with committed spend discounts use other cloud providers.[640]

5.201 Two hyperscalers submitted that the market research by Context Consulting commissioned by Ofcom[641] finds high levels of multi-clouding and switching, and low levels of concern around commercial practices and contractual issues.[642] [✂] submitted findings from a survey by Public First that cloud customers view committed spend discounts positively.[643]

5.202 In our view, these submissions cannot prove that committed spend discounts do not act as a barrier to multi-cloud. While many large customers do use multiple cloud providers for various reasons, this does not prove that there is no incentive to concentrated spend with a primary provider. The hyperscalers' analyses assume that any use of another cloud provider by their customers represents substantial multi-clouding. However, our evidence suggests that where customers do multi-cloud, their spend is generally concentrated around a primary provider and that customers usually procure only a small number of niche services from secondary providers (See Section 6). Moreover, we note that smaller cloud providers, such as Oracle and IBM, account for less than 5% of cloud infrastructure services sales in the UK – further suggesting that there is no material multi-clouding occurring between the hyperscalers and smaller cloud providers.

5.203 In addition, we do not consider that the survey evidence is relevant to our assessment of any impacts of cross-service committed spend discounts on multi-clouding. We note that these surveys were performed using broad samples of cloud customers in the UK, meaning they are unlikely to reflect the views of the small group of large and experienced cloud customers

---

[638] Hyperscalers also point to the findings of the survey of customers by Context, and the levels of customer multi-clouding among its respondents. The sample of customers used in this survey, being mid-sized enterprise customers, is likely to include few customers with individually negotiated committed spend discounts. [✂]; and Microsoft response to interim report, paragraph 162.

[639] The hyperscaler, [✂], also submitted estimates of the market shares of hyperscalers within the cloud market, estimating the three hyperscalers to have a combined share of roughly [✂]%. These estimates of the hyperscalers market shares and of the growth rates of other cloud providers differ substantially from our own, primarily due to [✂] inclusion of SaaS revenues in its figures. As discussed in Section 4, paragraphs 4.111 to 4.112, SaaS revenues are significantly more fragmented than those of IaaS or PaaS. [✂].

[640] This submission is based on the hyperscaler's analysis of publicly available information. [✂] confidential response to the interim report, [✂].

[641] This market research and what it can tell us about multi-clouding and switching is discussed in more detail in Annex 3.

[642] Microsoft response to the interim report, paragraph 168; and [✂].

[643] [✂].

which are eligible for individually negotiated committed spend discounts. Moreover, the survey did not ask customers about cross-service committed spend discounts specifically, and rather asked about committed spend discounts more generally. We do not find it surprising that customers reported that discounts are important for them in response to these questions, particularly given the depth that some of these discounts can take. As such, we do not consider that these surveys are necessarily informative about the individually negotiated committed spend discounts this section focuses on.

*Several customers told us that committed spend discounts represent an incentive to concentrate their spending with their primary provider*

5.204 We have gathered evidence which indicates that committed spend discounts have a material effect on large customers' decision making and contribute to a tendency among these customers to concentrate all or most of their spend with a primary provider.

5.205 In response to our interim report, we heard from several more customers whose decision making has been affected in various ways by their committed spend discounts. Some customers submitted that committed spend discounts create an incentive for them to procure multiple services from the same provider, rather than allocating incremental workloads to other providers.[644] Another customer told us that if its existing workloads with a provider it was committed to were not producing sufficient organic growth to gain an increased discount, broadening its usage clearly represented another way to achieve this.[645]

5.206 Some stakeholders who responded to our interim report said that the hyperscalers' committed spend discounts directly represented a means of locking in customers in a way which is difficult and expensive for them to break away from, and that they represent a way to disincentivise multi-clouding by customers and prevent smaller providers from growing their market share.[646] One stakeholder specifically identified the public sector as being affected by this, viewing the UK government's commitments to hyperscalers representing a commercial barrier to public sector customers using multiple cloud providers.[647] Decisions to concentrate demand with a single provider are also impacted by the targeting of specific workloads by hyperscalers, which has been mentioned by some customers.[648] For example, one customer told us that negotiations over spending commitments with cloud providers sometimes involve reference to specific workstreams and rival providers, and the promise of an increased discount on existing workloads if an additional workload is incorporated into the commitment. This customer told us that spending commitments to providers would represent an issue if it decided to move a workload to an alternative provider during the commitment period.[649]

5.207 From the evidence we have seen from customers during the market study, our assessment is that committed spend discounts are having a substantial effect on the decisions of many

---

[644] For example, [✂] confidential response to the interim report, page [✂] and Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[645] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[646] [✂] confidential response to the interim report question [✂] and Priyank Chandra response to the interim report, question 8.3.
[647] [✂] confidential response to the interim report question [✂].
[648] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂] and Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[649] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].

large customers, by encouraging them to concentrate all or most of their cloud spending with a primary cloud provider.

## Impact on competition for incremental workloads

5.208 For customers who already have some of their workloads in the cloud, committed spend discount have the potential to affect their choice of provider for incremental workloads, including workloads which they are bringing to the cloud for the first time and entirely new workloads. A provider may offer a customer which places an additional workload with them an increased discount on all of their spend, including existing workloads that a customer has with them.[650] This will be considered by the customer alongside the quality of providers' offering for the workload itself. The possibility of a deeper discount on the existing workloads with the hyperscaler can be a material incentive for the customer to place new workloads with the provider.[651]

5.209 We think this is a particular concern where customers face barriers to switching their existing cloud use. Under these circumstances, the customer is in a weaker bargaining position, as it faces a cloud provider who has a degree of market power over that customers' existing cloud usage. When this customer renegotiates its contract, the cloud provider may require the customer to increase the amount of spend they commit not to lose (some of) their current discount across their entire spending with them. Should the customer face a substantial cost to switching some or all their existing cloud use with the provider, the prospect of losing a discount can create a strong incentive to purchase incremental workloads from their existing provider, even where these workloads would otherwise be contestable.

5.210 In our interim report we described the experiences of several customers who told us that they had faced pressure during renegotiation with a hyperscaler ([✂]) to increase the size of their spending commitments. These customers told us that they had faced some degree of lock-in to the hyperscaler's products at the time of renegotiation, meaning it would be difficult and expensive to switch some of the workloads they had with the hyperscaler to an alternative provider, even another hyperscaler.[652] A customer we spoke to at the time of the interim report told us that the need to meet its high commitment, which it had been pressured into accepting during renegotiation, had caused it to shift specific workloads away from other providers to the hyperscaler.[653]

5.211 In this subsection we focus on the impact of the structure of the discounts on competition for individual workloads. First we consider the submissions hyperscalers have made on this topic, then we discuss the evidence we have gathered from our engagement with customers and other stakeholders.

*Hyperscalers' accounts of how commitments are agreed do not address our concern that customers are incentivised to concentrate their spend*

---

[650] In the interim report, we discussed the effects of customers being close to "thresholds" in hyperscalers' discounting schedules, where a customer may access a higher discount by increasing their spending a small amount. Given our understanding of the process of negotiation between customers and hyperscalers, where hyperscalers are able to make customers individualised offers of commitments and discounts, this framework is unnecessary to consider the effects of the discounts' structures on customers' incentives.

[651] Note that the concern does not require that the hyperscaler pressure the customer to purchase a greater volume or value of cloud services in total than the customer would require.

[652] Ofcom, 2023. Interim report, paragraph 5.155 to 5.159.

[653] Ofcom, 2023. Interim report, paragraph 5.157.

5.212 In response to our interim report, hyperscalers said that customers are free to choose the level of spend they wish to commit to and do not face pressure to increase their commitments at the point of renegotiation. Hyperscalers reiterated that customers choose their levels of commitment during renegotiations, and receive an offer of a discount based on the level of that commitment.[654] One hyperscaler told us that customers' use of cloud services tends to grow over time, and that customers generally seek a better discount upon renegotiation. This means that they generally want to commit their actual usage and forecasted growth to gain the biggest discount.[655]

5.213 In our view, these submissions by the hyperscalers are not consistent with the evidence we have received from customers, which indicates that there is considerable negotiation over the size of the commitment when customer renew their contract. Additionally, two hyperscalers acknowledged that they each work with customers to create a demand plan based on estimates of the projected usage of existing workloads and planned projects, and that these form the basis for the commitments customers are willing to make.[656] Even the few customers we spoke with which successfully resisted pressure to grow their commitments (at least to some extent) needed to push back and make counteroffers on the sizes of their commitments to the hyperscalers that they were negotiating with.[657]

5.214 A hyperscaler submitted that relatively few customers spend beneath their commitments, and that this indicates that their policies are not causing customers to set commitments which are too high relative to their actual spending.[658] In our view, this does not prove that commitments are not being set at a level which is affecting the purchasing decisions of customers. Since commitments function as minimum spends, customers which are not on course to reach their commitments have a strong incentive to take what steps they can to reach the commitment.[659] In any case, our concern is not that the structure of committed spend discounts will increase the total amount of cloud services consumed by customers, but that they will induce customers to purchase their cloud needs from a primary hyperscaler provider rather than using alternatives where appropriate.

5.215 In its response to the interim report, one hyperscaler [✂] submitted that customers with commitments generally spend more than their commitment with the provider they are committed to. It submitted that this provides customers with full flexibility to move any given workload they wish to an alternative provider during the duration of an agreement, because the commitment relates to overall spend rather than any specific workloads.[660]

5.216 We note that customers have the incentive to make their spending commitment as close to their anticipated spend as possible in order to secure the maximum available overall

---

[654] [✂] and Microsoft response to the interim report, paragraph 156.
[655] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[656] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂] and Microsoft response to the interim report, paragraph 159(ii).
[657] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]. Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[658] The hyperscaler submitted that customers may fail to reach their commitments for reasons such as lower-than-expected demand, and that in such cases they would seek to work with the customer and understand the underlying reason. [✂].
[659] A customer told us that they took steps such as shifting spending from one time period to another in order to reach commitments which they were not on course to reach organically within the available time. [✂] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[660] [✂].

discount, as acknowledged by this hyperscaler in its submission,[661] and that this includes renegotiating in-progress agreements in order to include any uncommitted spend in their commitment.[662] In our view, this means that the option of committing an incremental workload and receiving a higher discount is always available to customers which have an active committed spend discount, and as such will be a factor in their decision on which provider to use.

*Hyperscalers' analysis did not contradict the accounts we had from customers in relation to the pressures they face to increase commitments upon renegotiation*

5.217 One hyperscaler submitted that customers are not required to increase their commitments to receive the same level of discount, [663] suggesting that it does not pressure renewing customers into increasing their commitment. In support of this, this hyperscaler submitted an analysis showing that renegotiating customers who change their commitments by only small amounts see an insignificant change in the size of their discount.[664] In our view, this finding of statistical insignificance is in large part due to the small number of observations of customers who are defined as making small changes in their commitment size and it would be inappropriate to draw strong inferences from it. While the numbers of customers may be too small to draw robust statistical inferences, we note that [✂]% of [✂] customers which saw their discounts fall or remain stable following renegotiation actually increased the size of their commitment.[665]

5.218 Moreover, this analysis does not disprove the accounts of customers that they were pressured into increasing the size of their commitments during renegotiation. In fact, the hyperscaler's own data shows that most of its customers increase their commitments when they renegotiate their contract. In any case, we do not need to observe an increase in spending commitments when customers renegotiate for there to be a concern. As we explained above, the concern is that the structure of the discounts will incentivise customers to concentrate their demand around a primary provider. In other words, renegotiating customers who did not increase their commitment may simply have had no incremental workloads to bring to the cloud, while the committed spend discounts may still have discouraged them from moving some of their existing workloads to a rival cloud. As the committed spend discounts included in the hyperscaler's dataset only capture the outcomes of the negotiation process, it cannot reveal the process of negotiation. In particular, the hyperscaler's analysis does not capture the set of options available to customers during the negotiation process, which led them to choose these outcomes. The evidence gathered from customers on their experiences of negotiating with hyperscalers is a better source for understanding that process of negotiation.

5.219 A hyperscaler, [✂], submitted analysis of the levels of actual spend and committed spend immediately before and after the renegotiation of committed spend discounts. It submitted that customers' level of commitment typically increases at the point of renegotiation, whereas customers' actual spend does not immediately increase as a result. It submitted

[661] [✂].
[662] [✂].
[663] [✂].
[664] [✂].
[665] Of the [✂] customers who saw discounts fall following renegotiation, [✂] increased their commitments. Among the [✂] customers with the same discount following renegotiation, [✂] increased their commitments. Ofcom analysis based on [✂].

that this indicates that customers are choosing higher commitments upon renegotiation because their actual spend exceeded their previous commitment, rather than increasing their commitment because they have an incentive to do so, and that the lack of a sudden increase in spend following renegotiation indicates that customers are not increasing spend to match their commitments.[666]

5.220 In our view, the fact that customers typically increase their commitments to encompass a large portion of their previous uncommitted spend is consistent with the structure of the discounts encouraging them to maximise the size of their commitment. Indeed, the hyperscaler submitted that "there is a near one-to-one relationship between the average monthly actual spend prior to renegotiation and the average monthly committed spend in the re-negotiated contract." [667] As for the lack of sudden increases immediately following renegotiation, this is not something which we would expect to see if the commitment structure represented an incentive for customers to favour a primary cloud provider. The movement of workloads into the cloud and between cloud providers is unlikely to be instantaneous, and customers who increase their commitment on contract renewal are likely to increase their actual spend gradually over time. In this case, we would expect customers to meet their commitments towards the end of the commitment period rather than immediately after contract renewal.

*Several customers told us that they face pressure to increase commitments during renegotiation*

5.221 The customer accounts we detailed in the interim report described various methods that hyperscalers have used to pressure them to increase their commitments. Customers told us of: an expectation of a baseline level of growth at the point of renegotiation; the use of growth forecasts produced by hyperscalers based on past spend; the leveraging of wider support and benefits in negotiation over commitments, and; the use of information on customers' businesses such as pipeline projects and their reliance on provider-specific products.[668] Following our interim report, we engaged with additional customers and have learned more details about the mechanics of negotiation between customers and hyperscalers.

5.222 A number of customers told us that [✂] had proposed very large increases in the size of their commitment during renegotiation.[669]

5.223 Several customers told us that that while negotiating over the size of their commitments, they were told that they were expected to grow their spending.[670] One customer told us that this growth expectation was set at 20% per year.[671] Another customer told us that whilst there were options to save money by committing to a certain level of usage of specific services (using Reserved Instances or similar schemes) it was generally necessary to commit

[666] [✂].
[667] [✂].
[668] Ofcom, 2023. Interim report, paragraph 5.155 to 5.159.
[669] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]; Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]; Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]; and Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[670] Including Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂] and Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[671] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].

to growing their usage on an annualised basis in order to access the more comprehensive cross-service discount programmes. [672]

5.224 A number of customers told us that they understood that if they failed to grow their commitment at a sufficiently fast rate during renegotiation, they would face penalties to their discount level.[673] A customer even described being told directly during negotiations that it cannot expect to maintain its existing level of discount without a material level of growth in their spend.[674] Another customer told us that the possibility of failing to reach a new agreement and so reverting to paying list prices is a concern for it.[675]

5.225 Some customers described a process where [✂] would bring in people to audit their usage, and then present the customer with a commitment offer which offers the greatest discount to the customer for committing all of the possible spend they have identified.[676] [677] One customer said that its hyperscaler provider may hear about possible future projects it is planning, and incorporate them into its forecast and therefore commitment, despite these projects possibly not coming to fruition.[678]

5.226 Customers have described a high degree of inflexibility in the position of hyperscalers while negotiating committed spend discounts. One customer told us that despite spending considerable time and effort on negotiations, it was only able to slightly bring down the commitment demanded by its provider in its renegotiated agreement.[679] Another customer said that it spends significant resources on creating forecasts internally to prove or disprove the growth forecasts of its hyperscaler provider, and that only small changes in the size of the discounts it receives are possible in any case.[680]

5.227 The need to maintain consistent growth may be particularly pressing for customers who have reached a degree of maturity in their businesses and use of cloud. Such customers may be more likely to find their decisions on which providers to use for their cloud usage to be restricted by the need to reach their commitment, as they cannot rely on the organic growth of their business or the moving of new workloads to the cloud to grow their spending with their primary cloud provider.[681] However one customer, which was concerned about the situation that mature customers would find themselves in with respect to their commitments, told us that it believed that as the number of mature customers in the cloud market increases, it will become necessary for providers to begin to accommodate them.[682]

5.228 More generally, customers spoke of a lack of transparency during the negotiating process, with it being unclear how a change in the size or duration of their commitment would affect

---

[672] This customer also told us that the expectation for growth was relatively standard. Ofcom / [✂] meeting, [✂].
[673] For example, Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[674] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[675] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[676] For example, Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[677] We are aware that [✂] estimates a metric of [✂] for some customers, to identify the portion of a company's IT budget that is, or could be, allocated specifically to cloud-based services. The hyperscaler submitted that they estimate this figure using [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[678] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[679] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[680] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[681] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[682] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].

the size of their discount. Customers also told us that the negotiation process of discussing terms with [✂], and of waiting for them to make updated offers in response to changes in desired commitment or duration, can be very time-consuming. They have found themselves in the position of facing the expiration of their agreement and discounts, despite beginning the negotiations for their following agreement in good time.[683]

5.229 Some customers we spoke with have successfully pushed back on pressure to grow their commitments, including by threatening to switch to an alternative provider. Two customers told us that they had experienced pressure to grow their commitment on renegotiation but had managed to overcome this, to some extent, during their negotiations. Gaining this bargaining power is not necessarily easy or accessible to all customers. One of these customers re-architected its product at a significant cost to make it cloud agnostic, with the primary aim of gaining bargaining power over its provider.[684] The other customer recognised that a customer that lacked its size may have more difficulty in pursuing a "conservative approach" in negotiations.[685]

5.230 In summary, the customer evidence we have gathered over the course of the market study shows a consistent pattern of [✂] using commercial pressure to influence its customers to increase their commitments over subsequent renegotiations. For some customers, the level of growth in commitment is not problematic as they are growing their cloud usage fast, or because they have bargaining power, they can use to mitigate the pressure. However, for other customers, the expectation of growth in commitments may create strong pressure to move workloads to [✂] and represents a major barrier to them using other cloud providers. Our view remains that this conduct can be a barrier to multi-cloud and has the potential to impact competition for customers' incremental workloads.[686]

# Potential impact on competition

## Impact on smaller cloud providers

5.231 Customers with committed spend discounts are relatively large and generally have multiple cloud workloads and require a large number of distinct cloud products.[687] Hyperscalers, with their broad product ranges across the cloud stack, will generally be able to offer such customers a solution for all or most of their cloud needs. Smaller providers, however, have narrower product ranges and so may only be able to compete for some workloads of a customer's overall cloud needs.

5.232 If a customer chooses to purchase one or a few workloads from a smaller provider, the customer may be forgoing an increase in their discount on the workloads that customer must purchase from a hyperscaler. In that case, the smaller provider must be able to compensate the customer for this loss in discount to win the individual workload. We are concerned that this reduces the opportunity for smaller cloud providers to compete for components of the demand of these large customers.

---

[683] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂] and Ofcom/ [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[684] This customer noted that bargaining power can also be used to gain benefits besides an increased percentage discount during negotiations, such as cloud credits. Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[685] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[686] Ofcom, 2023. Interim report, paragraph 5.160.
[687] See Section 5 "Technical barriers" for further details.

5.233 While customers with spending commitments to a hyperscaler are only a small proportion of the customers of cloud services by count, they account for a large proportion, [✂]%, of hyperscalers' total cloud revenues.[688] Given the prevalence and strength of economies of scale in cloud infrastructure services, barriers to smaller providers winning some of the demand of large customers has the potential to inhibit their ability to gain scale – raising barriers to entry and expansion. This concern is particularly relevant to the market leaders as they have the largest established customer base.

5.234 The multi-year durations of the agreements which include spending commitments may compound their impact on the wider market. The prevalence of agreements with long durations means the amount of cloud spending at any given time that rival providers are able to compete for and win is restricted.[689] Additionally, one consultant to cloud customers expressed a view that cloud providers prefer to get customers locked in to longer-duration contracts and that this is inadvisable for customers unless they have a robust negotiating team in place allowing for appropriate contract amendments to be made.[690]

5.235 In our interim report, we gave the views of some smaller cloud providers on the effects that the hyperscalers' committed spend discounts have on competition in the cloud market.[691] One provider told us that the structure of the discounts induces customers to move as many workloads to the hyperscaler as possible and causes customers to sometimes choose sub-par solutions to meet commitments. This provider also said that the mechanisms of the hyperscalers' discounts which lead to this are not necessary to achieve the legitimate goals of customer spending commitments.[692]

5.236 Smaller cloud providers we spoke with after the publication of our interim report told us that committed spend discounts can have positive effects for customers and the functioning of the cloud market, but have the potential to be used in anticompetitive ways. This includes by limiting customers' ability to switch, by worsening lock in, and by creating an incentive to single source.[693] One smaller provider told us that committed spend discounts represented one barrier among several that the hyperscalers have created to them winning the business of larger customers, and that due to the size and importance of this customer group within the cloud market, this represents a major barrier to their growth generally.[694]

5.237 Similarly, an intermediary told us that the hyperscalers' use of committed spend discounts creates a price barrier for other providers, meaning that customers would face increased prices from using a competitor even on a trial basis, and ultimately preventing entrant providers from reaching a viable scale.[695]

---

[688] [✂] response dated [✂] to follow-up questions dated [✂], concerning s.174 notice [✂], question [✂]; [✂] response dated [✂] to follow-up questions dated [✂] concerning s.174 notice [✂], question [✂]; and [✂] response dated [✂] to follow-up questions dated [✂] concerning s.174 notice [✂] question [✂].
[689] Ofcom, 2023. Interim report, paragraph 5.136(d).
[690] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[691] Ofcom, 2023. Interim report, paragraph 5.147.
[692] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[693] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂] and Ofcom / [✂] meeting, [✂].
[694] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[695] Ofcom / [✂] meeting, [✂].

## Impact on ISVs

5.238 In our interim report we noted that the effects of committed spend discounts on ISVs will be different to other smaller cloud providers, as spending on hyperscalers' marketplaces will draw down customers' commitments.[696] ISVs told us that committed spend discounts still have the potential to affect them in the same ways as other smaller cloud providers, including due to the hyperscalers' policies which do not treat marketplace spend with parity to spending on hyperscalers' own products.[697] Our view from the evidence was that the ability to draw down their commitments with marketplace spending represents a major reason for customers to use hyperscalers' marketplaces.[698]

5.239 One hyperscaler, in its response to our interim report, said that the eligibility of marketplace spending to draw down customers' commitments means that ISVs whose products are available on its marketplace are not in the position of needing to use their own committed spend discounts to win the business of customers who have spending commitments.[699]

5.240 Another hyperscaler submitted that its policy allowing marketplace spend to draw down customers' commitments was positive for ISVs at it promotes their offerings.[700] A smaller cloud provider who responded to our interim report submitted that it views the eligibility of marketplace spending to draw down commitments as a powerful generator of network effects for the hyperscalers, and a way for hyperscalers to incentivise customers' reliance on a single cloud provider.[701]

5.241 The effects committed spend discounts have on ISVs appear mixed. Overall, the evidence continues to support the view that committed spend discounts are an important driver of customers' use of hyperscalers' marketplaces. Ultimately, the consequences of this for ISVs and the cloud market will depend on the policies of hyperscalers' marketplaces and the degree to which they are managed in a way which is fair to ISVs and their products, especially in the longer term.

## Impact on competition between hyperscalers

5.242 In our interim report, we reported that each of the hyperscalers emphasised its view that they compete to win customers' commitments as a normal part of competition within the market.[702] There are circumstances where this is likely to be true. The hyperscalers, with their wide ranges of products across the cloud stack, are likely to be able to compete on an equal footing for a customer's spend commitments when they first migrate to the cloud, or where customers can easily switch their existing workloads between cloud providers when renegotiating contracts.

5.243 However, we also noted that sometimes a cloud provider may have a certain degree of market power over workloads for which a customer is less able to choose an alternative provider. In such cases, hyperscalers may find themselves in a similar position to smaller providers, attempting to win a subset of a rival customer's workloads instead of being able to compete for all or most of that customer's needs. For example, where a customer has

---

[696] In some cases, a cap applies to the proportion of a customer's total commitment which may be drawn down through marketplace spending. Ofcom, 2023. Interim report, paragraph 5.149.
[697] Ofcom, 2023. Interim report, paragraph 5.151.
[698] Ofcom, 2023. Interim report, paragraph 5.150.
[699] Microsoft response to the interim report, paragraph 159(v).
[700] [✂].
[701] [✂] confidential response to the interim report, page [✂].
[702] Ofcom, 2023. Interim report, paragraph 5.143.

specific needs which are only met by a given hyperscaler, rival hyperscalers may not be able to offer the full range of products that match the customer's needs. Existing customers may also face high switching costs for their existing workloads with their current cloud provider, meaning that a rival hyperscaler can only compete for incremental (and sufficiently siloed) workloads of that customer.

5.244 In these cases, a hyperscaler may be able to use the workloads over which it has a degree of market power to incentivise its customer to concentrate all of its needs with that hyperscaler. For example, where a customer faces a substantial cost to switching some or all their existing cloud use, the threat of losing a discount on these existing workloads could create a strong incentive to purchase its incremental requirements from the ecosystem of their existing provider. This may make it harder for rivals to compete for incremental workloads which would otherwise be contestable, including other hyperscalers that offer a broad range of cloud services.

## Hyperscalers' rationales for their use of committed spend discounts

### Each of the hyperscalers submitted that discounts are a pro-competitive feature of the cloud market

5.245 In our interim report we noted that discounts can provide lower prices to the customers receiving them and be a basis for competition between providers.[703] In response, each of the hyperscalers submitted that discounts are a positive feature of markets, allowing large customers to exercise their bargaining power, indicating strong competition between providers and directly providing customers with lower prices.[704] In our view, there are many forms that discounts can take, and cloud providers do offer a wide range of different discount types with different structures. Our concern does not relate to the practice of price discounting generally, but to the specific structure of the committed spend discounts used by hyperscalers.[705]

5.246 One hyperscaler submitted that committed spend discounts cannot encourage customers to purchase a "packaged business application" of cloud services from a single hyperscaler's ecosystem because the cross-service nature of the commitment and the discount means that customers are free to draw down their commitment using any mix of services, including a single service if they choose, meaning no "package" is necessary.[706] In our view, the agnostic nature of the discount to the exact services purchased does not undermine the ability for committed spend discounts to encourage customers to purchase all or most of their needs from a single cloud provider. This is because allocating any workloads to a

[703] Ofcom, 2023. Interim report, paragraph 8.46.

[704] Hyperscalers also submitted that committed spend discounts allow efficiencies in investment due to improved predictability in customer spending, which are passed on to customers in the form of lower prices. These submissions are considered later in this section. [✂]; [✂]. [✂]. [✂] confidential response to the interim report, paragraph [✂].

[705] One hyperscaler submitted that committed spend discounts allows cloud providers to compensate customers for the cost of switching and migration and so promotes competition between hyperscalers. In our view, it is not clear how this function of discounting requires the specific structure and features of the hyperscalers' cross-service committed spend discounts, outside of the narrow case of competing with a directly equivalent committed spend discount. [✂].

[706] [✂].

secondary provider risks them losing out on a deeper discount across their entire usage with their primary provider.

## Hyperscalers state that commitments encourage investment

5.247 In our interim report we highlighted a rationale hyperscalers have stated for their use of committed spend discounts. They said that customer commitments allow cloud providers to forecast the future demand for their services more accurately and with more certainty, meaning providers are able to invest more confidently.[707] One hyperscaler submitted that the benefits from these improved forecasts are passed on to customers through the discounts they receive from making commitments.[708] In response to our interim report, hyperscalers reiterated this view. They submitted that cost savings as a result of increased efficiency are passed on to the customers who are making commitments.[709]

5.248 One hyperscaler said that committed spend discounts enable the launch of new services, ensure that they have the appropriate capacity for existing services, and allow providers to allocate their resources across the range of services they offer.[710] In our view, these reasons do not seem to apply clearly to the cross-service nature of the discounts. Customers may choose to use any of the hyperscalers' wide range of products to draw down its commitment, meaning they do not give providers certainty about the usage of specific products.[711]

5.249 While it may be the case that greater certainty of future revenues aids investment, this does not justify many of the particular characteristics of the committed spend discounts used by the hyperscalers. In particular, the degree of individualisation in the size of discount received by different customers for similar commitments does not appear to support the hyperscalers' suggestion that these discounts reflect the pass-through of cost savings. Additionally, the long durations of the commitments means that customers use of cloud services could follow any pattern during the contract period, including large peaks in usage, meaning they do not give certainty to providers of the overall capacity needed to accommodate that customer.

5.250 Other ways of structuring commitments and discounts are available to hyperscalers which can provide this certainty without leading to the concerns we have identified. For example, spending commitments relating to individual services only.[712]

## Hyperscalers submitted that commitments provide customers with greater certainty

5.251 Hyperscalers submitted that making a spending commitment allows customers to overcome difficulties in predicting their cloud spending and give greater price transparency, because of

---

[707] In response to our interim report, one hyperscaler submitted that the certainty created by spending commitments also allows for greater customer-specific investments. Our assessment of these submissions is also applicable to this. Microsoft response to the interim report, paragraph 165.
[708] Ofcom, 2023. Interim report, paragraph 5.144.
[709] [✂]; [✂]; and [✂].
[710] The hyperscaler also submitted references to several academic papers on firms' incentives to invest under uncertainty and on patents, which show a mixture of findings. [✂].
[711] The hyperscaler submitted that the same rationale does apply to cross-service commitments, but did not articulate this view further. [✂].
[712] A cloud provider told us that they are able to gain sufficient certainty of use to support their investment through their committed spend discounts, which are of a shorter duration and which are workload-specific. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

the tools they provide on ongoing usage and spending.[713] In our assessment, the availability of such monitoring tools is not dependent on committed spend discounts. Committed spend discounts also do not prevent customers from unexpectedly spending above their commitments, only from unexpectedly spending below them.

### Smaller providers' use of committed spend discounts do not cause the same concern

5.252 Various types of discounts which involve customer commitments are used by smaller providers. Some of these discounts are individually negotiated with customers and so resemble the committed spend discounts of the hyperscalers in structure.

5.253 One cloud provider submitted that when smaller cloud providers compete with hyperscalers, using committed spend discounts allows them to compete more effectively and win business.[714] A hyperscaler submitted that smaller providers can use committed spend discounts to reduce the uncertainty around future demand associated with entry, and support their investment.[715]

5.254 One cloud provider told us that the structure of its committed spend discounts differs from those of the hyperscalers. This provider explained that its contracts are shorter in duration than those of the hyperscalers and that they apply only to specific workloads, both in terms of the eligibility of spending to draw down the commitment and the application of the percentage discount.[716] Another smaller cloud provider submitted that it offered discounts for commitments, though primarily for specific IaaS products rather than across its product range.[717]

5.255 Given the narrower product ranges of smaller providers, in principle, their use of such discount structures is less likely to limit the ability of rivals to compete for some of their customers' demand. Our concern is with the structure and use of committed spend discounts by the hyperscalers, not with those of smaller providers that have a narrower range.

## Conclusion on committed spend discounts

5.256 We remain of the view that price discounting can be a means of competition between cloud providers and so has the potential to benefit customers and lead to lower prices. We also accept that it can provide greater certainty to providers and customers, encouraging investment in product innovation.

5.257 However, we are concerned that the structure of the hyperscalers' committed spend discounts has the potential to encourage large customers to use a single hyperscaler for all or most of their cloud needs. In particular, customers can receive a larger discount on their entire spend with a provider by maximising the workloads that they place with a single provider. The individualised and private nature of these discounts means that hyperscalers can tailor these discounts in a way that encourages customers to place most of their workloads with their primary provider. Overall, we therefore consider that committed spend discounts can pose a barrier to multi-cloud.

---

[713] Microsoft response to the interim report, paragraph 161. [✂].
[714] [✂].
[715] [✂].
[716] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[717] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

5.258 For smaller cloud providers that do not offer the full range of products, the use of committed spend discounts by the hyperscalers may further raise barriers to entry and expansion. Due to their narrower set of products, they may find it difficult to compete for a customer's full commitment spend as they may be unable to match the hyperscalers' discounts across a wide set of a customer's product needs. They are more likely to rely on being the secondary provider of a customer and the use of committed spend discounts by hyperscalers may limit their ability to compete for a subset of customers' workloads that otherwise may have been contestable.

5.259 In principle, the hyperscalers have sufficiently wide product portfolios to compete for the full commitment of most large customers when they first move into the cloud. This could encourage competition between the hyperscalers by strengthening the bargaining position of these customers. However, where customers have a limited ability to choose an alternative cloud provider for some of their needs, even hyperscalers will only be able to compete for some workloads of these large customers. In these cases, committed spend discounts could make it more difficult for larger rivals to challenge for this subset of workloads.

# Challenges predicting cloud spend

5.260 This subsection sets out our assessment of whether difficulties forecasting usage, interpreting prices and, consequently, predicting cloud spend may create additional barriers for customers wishing to switch and/or multi-cloud.

5.261 Predicting future cloud spend is a function of forecasting future usage, and knowing what price services will be charged at. Below we set out the evidence on each of these elements in order: forecasting usage, interpreting prices, and predicting cloud spend. The subsection concludes with an assessment of whether this creates additional barriers to switching and multi-cloud.

## Forecasting usage

5.262 The cloud services market is complex and dynamic, and the purchasing of cloud solutions involves many parameters. The complexity and pace of marker developments can make it difficult for customers to forecast their future needs. For example, Microsoft Azure launched with a handful of services and now has over 200, similarly AWS started with three services in 2006 and has over 240 services today.[718]

5.263 A cloud provider, [✂], noted that when customers begin their migration to the cloud, they sometimes find it difficult to anticipate their future demand, but recognised that this can be because customers find it difficult to anticipate how their cloud needs will evolve.[719] We also heard that many customers lack the required skills and knowledge to understand the long-term implications of their choices.[720] We have heard that customers tend to forecast on a year-on-year basis and are unlikely to be able to forecast accurately much beyond this.[721]

[718] [✂].
[719] [✂].
[720] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[721] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]; Ofcom / [✂] meeting, [✂]; and Ofcom / [✂] meeting, [✂].

5.264 Given the length of time that it may take enterprises to migrate workloads to the cloud, and the potential for customer needs to change as new use-cases arise, their businesses evolve, and new products become available, it is impossible for customers to accurately plan future demand at the initial migration point.

## Interpreting prices

5.265 In our market research, 52% of customers said they were concerned or very concerned about the lack of pricing transparency.[722] Customers in the qualitative part of the market research told us that 'bills can be confusing'[723] and they find that cloud providers are often 'trying to muddy the waters in terms of costing'.[724]

5.266 Other customers had a similar experience. BT Group's CFI response highlighted the concern that a lack of transparency of egress fees can adversely affect market dynamics both now and in the future.[725] [✂] also told us that they find it very difficult to forecast cloud spend in part due to a lack of price transparency beyond trial periods and over potential cost savings.[726] Some customers [✂] also told us that hyperscaler pricing can be difficult to understand.[727] One cloud provider [✂] also told us that "this difficulty in reading hyperscalers prices has a pernicious effect on customers" as they are not "guaranteed the best price/performance ratio when they sign up to a cloud provider and often face 'hidden costs' due to the complexity".[728]

5.267 For some customers this lack of pricing transparency when they first move into the cloud can make it difficult for them to understand the reality of their contracts. In some cases, we have heard examples of customers experiencing 'bill shock' which can be due to a lack of clarity and transparency, and challenges understanding what they are paying for, how much they are paying and why.[729]

5.268 [✂] told us that its challenges with forecasting cloud spend are not based around transparency of pricing, which it thinks is high. Instead, it said that different pricing models can make it difficult to understand what costs will look like in the future.[730] This point was echoed in responses to our interim report, where some stakeholders called for simplification of pricing models to help resolve issues around pricing transparency.[731]

## Predicting cloud spend

5.269 Predicting what a customer's future cloud spend would be is a function of their future needs, and the prices providers charge for those services. Findings from our market research show that in one in three cases, IaaS/PaaS customers find it difficult to accurately predict the

[722] Context Consulting research report, slide 131.
[723] Context Consulting research report, slide 33.
[724] Context Consulting research report, slide 58.
[725] BT Group response to the CFI, page 25.
[726] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[727] Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂], Ofcom / [✂] meeting, [✂].
[728] [✂] confidential response to the interim report, page [✂].
[729] Ofcom / [✂] meeting, [✂]; Priyank Chadra response to the interim report, page 11; and TechUK response to the interim report, page 2.
[730] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[731] [✂] confidential response to the interim report, page [✂]; Federation of Communication Services response to the interim report, page 6; Name withheld 1 response to the interim report, page 7.

future costs of their cloud computing.[732] This is echoed by some of the evidence outlined above on forecasting usage, and interpreting prices.

5.270 However, these are not issues that every customer faces. We heard from [✂] and [✂] that they find forecasting cloud infrastructure spending relatively issue-free,[733] and they do not encounter 'high volatility of prices or find pricing particularly opaque'.[734]

5.271 We noted in our interim report that we have seen some evidence of cloud providers awareness of the issues with transparency of billing. [✂] documents showed that customers have told it that it is becoming increasingly complex to compare the true cost of [✂] with competitors. It noted that this complexity comes from different discount plans and pricing structures.[735]

5.272 In response to our interim report, hyperscalers stated that prices are listed publicly on their website, which are available for all customers to use.[736] For some customers hyperscalers also pointed to committed spend discounts as a way for customers to gain predictability in their cloud spend.[737] However, we note that committed spend discounts only offer customers certainty on their minimum spend and do not prevent customers from unexpectedly spending above their commitments.[738]

5.273 Our interim report also explored cost monitoring services that are available from hyperscalers.[739] We have heard some positive commentary on the services available from hyperscalers,[740] including customers in the qualitative part of the market research who told us how they had seen improvements in cost monitoring services available from cloud providers.[741] Other customers have pointed out their limitations, including a lack of accuracy when forecasting.[742] Our market research showed a consensus among customers that vendors could do more to help manage their cloud spend.[743]

5.274 We have also heard evidence from some smaller cloud providers about the challenges customers face when trying to forecast spend.[744]

a) Hyve Managed Hosting told us that:[745]

i) Customers often do not know the real cost of the cloud services they plan to consume until they have started using them and receive their bills. When describing

---

[732] Context Consulting research report, slide 95.
[733] [✂] response dated [✂] to the s.174 noticed dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[734] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[735] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[736] AWS response to the interim report, page 13, paragraph 29; Microsoft response to the interim report, page 31, paragraph 116; and [✂].
[737] Microsoft response to the interim report, page 9, paragraph 28 and Google response to the interim report, page 3, paragraph 14.
[738] For example, we find that around [✂]% of [✂] first deal customers with completed contracts had an actual spend that was twice the amount of their commitment. Ofcom analysis based on [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[739] AWS Cost Explorer, Azure Cost Management and GCP Billing.
[740] Ofcom / [✂] meeting, [✂], Ofcom / [✂] meeting, [✂].
[741] Context Consulting research report, slide 80.
[742] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂] and Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂].
[743] Context Consulting research report, slide 98.
[744] [✂] confidential response to the interim report, page [✂].
[745] Ofcom / CISPE / Hyve Managed Hosting meeting, 2 February 2023.

the problem, it highlighted the complexity of the issue and noted a series of examples.

ii) Hyve Managed Hosting also explained how some cloud providers bundle their products, making it difficult for customers to compare costs between providers as everything is integrated into one price. This can be particularly prevalent when looking at certain software licensing. In other instances, it can be the complexity associated with individual components that make it difficult for customers to understand what they are paying for and predict future usage. For example, when building a quote for public cloud vendors, in order to gain an accurate cost a customer may need to understand highly complex requirements such as how many requests per second their load-balancer would be receiving, or how many disk reads they would be carrying out per second. These challenges also make it difficult for customers to compare prices between providers, especially when different providers measure different metrics.

b) [✂] also told us that the hyperscalers introduce unexpected additional charges as companies begin to scale their cloud usage, making it difficult for customers to forecast spend.[746]

## Conclusion on predicting cloud spend

5.275 Overall, the evidence indicates that some customers may find it difficult to accurately predict the future costs of their cloud computing. These challenges may stem from the inherent complexity or undefined nature of a customer's needs which makes it difficult to forecast future usage of cloud services. It may also be exacerbated by difficulty interpreting and comparing pricing across providers or a lack of clarity on the exact value of prices and fees.

5.276 Whilst we acknowledge that some customers have told us they do not face the challenges we have identified, where customers do find it difficult to predict and compare future spend when choosing a provider, this may prevent them from exercising effective choice in selecting the cloud architecture that is most appropriate for their needs. This may reduce the effectiveness of competition based on quality and price, which in turn compounds the barriers to entry and expansion for smaller cloud providers for whom price is a key lever to attract new customers. Difficulties predicting usage can also lead to "bill shock" for some customers, and make it hard for customers to protect themselves from potential future harm when they first migrate to the cloud (we detail this further in Section 8).

# Conclusion on barriers to multi-cloud and switching

5.277 We conclude that the barriers we identify in this section combine to substantially limit the ability of some customers to multi-cloud and to switch and so limit competition in cloud infrastructure services.

a) Technical barriers increase the costs to customers of adopting multi-cloud architectures and switching. This is particularly the case where more integrated forms of multi-cloud and switching between clouds involve a significant number of applications which are

[746] Ofcom / [✂] meeting, [✂].

tightly integrated with proprietary services of the cloud provider. Technical barriers are also likely to pose material limitations on some multi-vendor deployments.

b) Egress fees hinder the adoption of multi-cloud architectures which require substantial movement of data between workloads hosted on different clouds. They can also increase the costs of switching between cloud providers for those customers who would have to transfer large amounts of data and more generally as switching will typically involve more integrated forms of multi-cloud as an intermediate step.

c) Committed spend discounts encourage large customers to purchase all or most of their cloud needs from their primary cloud provider, reducing their incentives to multi-cloud.

5.278 Some customers will be less affected by these barriers. This may be true for customers who host limited amounts of data in the cloud and have relatively simple requirements which can be easily integrated and moved between cloud providers. However, our evidence suggests that a significant number of users have had to settle for a sub-optimal approach to multi-cloud and are concerned about their lack of ability to switch.

5.279 The barriers we have identified appear particularly material for customers with large and complex cloud architectures, customers who need to adhere to specific regulatory requirements and/or customers who are less technically sophisticated. These customers are likely to account for a substantial fraction of demand for cloud infrastructure services in the UK, at least in revenue terms. We also believe they encompass many critical sectors, such as government, financial services, healthcare, social media, as well as our core sectors of broadcasting and telecoms.

5.280 Because of these barriers to multi-cloud and switching, a customer's initial choice of cloud provider is important. Once a customer chooses a cloud provider, these barriers make it more likely that they will concentrate their usage within that cloud providers' ecosystem. They are also more likely to increase their usage from the chosen cloud provider's ecosystem as they migrate more workloads into the cloud and their needs evolve. This is particularly the case for those customers who face material barriers to multi-cloud and switching.

5.281 As a result of these dynamics, much of competition in cloud infrastructure is currently centred around attracting customers when they first migrate into the cloud. This is evident in cloud providers' strategies, with widespread practices of cloud providers offering new customers a range of incentives (such as free trials) to win their business. However, barriers to multi-cloud and switching have the potential to lessen effective competition in cloud infrastructure services:

a) Cloud providers, particularly AWS and Microsoft, face a weaker threat from some customers adopting a multi-cloud architecture or switching to make savings or purchase a rivals' best-in-breed solutions. These barriers also reduce the potential for workload competition to put pressure on these customers' primary cloud provider. This is particularly the case where a customer purchases a range of tightly integrated services from a cloud provider's broad ecosystem, and this can impact a customer's existing and incremental workloads.

b) Where smaller cloud providers are unable to compete for a customer's entire cloud needs, barriers to multi-cloud and switching can further raise barriers to entry and expansion which we discuss in Section 6. They make it more difficult for those providers to challenge for one or several of a customers' workloads, which could inhibit their ability to grow their customer base and gain scale.

5.282 Challenges predicting cloud spend can further limit competition by preventing customers from exercising effective choice in selecting the cloud architecture that is most appropriate for their needs. This can reduce competition between cloud providers based on quality and price.

# 6.Barriers to entry and expansion

6.1 In this section we identify market features which are important for competing in cloud infrastructure services and consider the extent to which these create barriers to entry and expansion.

6.2 Competing as a vertically integrated cloud provider requires significant and on-going investment in several different areas:

a) physical infrastructure, including data centres, servers and network equipment;
b) a broad product portfolio (including solutions for specific industries); and
c) customer acquisition strategies.

6.3 The investments made in all three of these areas have the potential to raise barriers to entry and expansion. There are also other potential barriers to entry and expansion, including the need to attract highly specialised teams of engineers and product developers.

6.4 ISVs may face lower barriers to entry and expansion to supply PaaS, as they typically use the infrastructure (IaaS services) of other cloud providers to provide their services rather than using their own infrastructure. However, the potential reliance of ISVs on the hyperscalers may raise other issues, which we discuss in Section 7.

## Investment in physical infrastructure

6.5 In this subsection, we consider the extent to which investments in the physical infrastructure required to provide cloud services may act as a barrier to entry or expansion.[747]

6.6 This subsection is structured as follows:

a) Investment in data centres;
b) Economies of scale in data centres;
c) Innovation in underlying hardware; and
d) Benefits from sharing cloud infrastructure with non-cloud businesses.

### Investment in data centres

6.7 The hyperscalers have hundreds of data centres located around the world – for example, Microsoft Azure operates [✂] [between 200-350], with plans to build more in the foreseeable future.[748, 749] In comparison, smaller providers operate fewer data centres – for example, IBM operates [✂] [between 50-100] and OVHcloud operates 37 data centres.[750]

[747] ISVs also rely on physical infrastructure to provide PaaS (and SaaS) services, but they typically use the infrastructure (IaaS services) of other cloud providers, to provide their services rather than using their own infrastructure.
[748] Microsoft response dated 9 December 2023 to the s.174 notice dated 21 October 2022, Part B question 26a (Confidential Annex B26).
[749] Microsoft News, 2021. Microsoft's virtual datacenter grounds 'the cloud' in reality [accessed 15 September 2023].
[750] IBM response dated 23 December 2023 to the s.174 notice dated 25 October 2022, Part B question 24a; OVHcloud response dated 12 December 2022 to the s.174 notice dated 27 October 2022, Part B question 17a. OVHcloud website. Datacentres [accessed 15 September 2023].

**Table 6.1: Data centre figures**[751]

| | Global data centre count (owned and leased / co-located, 2021) | Average global data centre capacity (MW per data centre, 2021) |
|---|---|---|
| **AWS** | [✂] | [✂] |
| **Google** | [✂] | [✂] |
| **Microsoft Azure** | [✂] | [✂] |
| **IBM** | [✂] | [✂] |
| **Oracle** | [✂] | [✂] |
| **OVHcloud** | [✂] | [✂] |

*Source: Ofcom analysis based on responses to s.174 notices, and data from provider websites where available.*

6.8 Costs of data centres are very high and can require significant capital expenditure on fixed assets such as data centre premises, servers and network equipment. Operating costs of data centres can also be high and include costs of labour, support services, and networks (e.g. energy, internet, etc.).[752] For example:

a) A co-authored 2016 study by Emerson Network Power and Ponemon Institute estimated that the average cost of a data centre is more than £3.6 million per year, and can amount to around £6 million per year for some data centres.[753]
b) A 2022 report by Dgtl Infra states that it could cost between $7-12 million per megawatt of commissioned IT load to build a data centre.[754] This means that build costs per data centre alone could be significant, in the tens to hundreds of millions.

6.9 As such, the need to invest in a network of data centres could act as a significant barrier to entry and expansion for cloud providers.

6.10 Barriers to entry may be reduced by leasing the property and/or co-locating in existing data centres, as this reduces the high capital costs associated with construction and other barriers such as the lead time of building data centres. There is evidence that small-scale

[751] Providers may be using some of their data centre capacity for purposes other than providing cloud services.
[752] Based on analysis of responses to our statutory information requests.
[753] The study was sponsored by Emerson Network Power and conducted by Ponemon Institute, which is a research centre focused on privacy, data protection and information security policy. Estimated costs in the study are based on data from 41 data centres, representing 31 companies, who reported their costs in the following cost categories: physical plant (amortized), IT assets (amortized), operating costs (including labour costs), and energy costs. Converted from 4.93 million USD and 8.10 million USD using the average exchange rate in 2016 of 1 USD to 0.74 GBP. Emerson & Ponemon Institute, 2016. Cost to Support Compute Capacity. Retrieved from Wayback Machine [accessed 29 September 2023].
[754] These estimates include the following cost categories: land and building shell, electrical systems, HVAC / mechanical/ cooling systems, building fit-out area. Dgtl Infra, 2022. How much does it cost to build a data centre? [accessed 19 September 2023].

public cloud providers have entered the market in recent years by buying co-location space from providers of data centres, e.g. Civo[755] and OVHcloud.[756] [✂].[757] [✂].[758]

6.11 Capital costs can still be substantial where providers lease or co-locate data centres due to investments in IT equipment (such as servers). For example, information provided by [✂] suggests that the majority of its data centre capital expenditure (capex), relates to investment in IT equipment (such as server components, server production and network switches), with construction costs being a significantly smaller share of its total capex.[759]

6.12 Regardless of whether a cloud provider is constructing or leasing, there can be capacity challenges that make it more difficult for cloud providers to acquire additional data centre space. While this can impact all cloud providers, this may particularly affect potential entrants who may find it difficult to replicate the offering of providers who have already secured some resources:

a) Capacity shortages in key locations – [✂] told us that some cloud customers (e.g. e-commerce, banking, or gaming companies) require the latency period for their information systems to be as low as possible, which implies that data centres need to be located in closer geographical proximity to companies' premises. In densely populated regions such as in Greater London, the available real estate to build new data centres is limited. As a result, [✂] suggested that limited access to real estate is an important barrier to entry and expansion.[760] Similarly, [✂] also said that capacity can vary by location. For example, it said there is currently available capacity in the UK generally, but capacity has become more constrained in west London, due to a lack of power capacity in certain key substations which has delayed deployment of new capacity. [✂] also highlighted challenges regarding scarcity of available development land suitable for data centres – but noted that these capacity issues apply to building as much as leasing/co-location.[761]

b) Availability challenges in co-location data centres – if there are shortages in co-location space in existing data centres in popular regions, smaller providers may find it more difficult to compete with the hyperscalers for co-location space. [✂] said that if there is

[755] Civo launched its cloud platform in October 2021. It rents colocation space from data centre owners and deploys its own servers. Civo has servers in the UK, Frankfurt, New York and Phoenix, and plan to launch further locations in India, Singapore, the US and the UK. Ofcom / Civo meeting, 18 October 2022.
[756] OVHcloud rented its first data centre in 2001, followed by a fully owned data centre two years later and today has 37 data centres in 8 countries around the globe. OVHcloud website, Datacentres [accessed 15 September 2023].
[757] [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂]; and [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂].
[758] [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to our s.174 notice dated [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂].
[759] [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂].
[760] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂].
[761] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂].

available capacity, it is relatively easy to rent additional co-location space, but things become more challenging if there are limitations in available space or power.[762] Similarly, [✂] noted that co-location data centre supply can fluctuate – capacity availability depends on when co-location providers expect to have capacity becoming available. Global supply chain issues have extended timelines too for co-location providers to build capacity, exacerbating availability challenges.[763]

6.13 In response to our interim report, one cloud provider [✂] stated that it is not necessary to build a large-scale operation on initial entry and therefore does not consider investment in physical infrastructure to be a significant barrier to entry. Instead, it [✂] suggested that small and mid-scale providers can enter the market by using a variety of capital efficient strategies (including leasing and co-location) and grow their capacity and geographic reach as their business expands. [764] We recognise that it is possible to start providing IaaS services with a single data centre and add data centres incrementally, as illustrated by the entry of both small and mid-scale cloud providers. We observe that small scale cloud providers have much fewer data centres and are located in fewer geographic regions.[765] Of the mid-scale cloud providers, Oracle has been able to make the most progress in this area. Oracle entered the market in 2016 with one public cloud data centre and has since added data centres at a fast rate (it now operates [✂] [between 150-200] data centres globally).[766] Oracle has also launched public cloud data centres in several geographic regions to rival AWS, Microsoft and Google.[767]

6.14 However, some customers value having access to a large network of data centres and may consider that only the hyperscalers' offerings fit these requirements.[768] It is important to offer cloud services in multiple geographic regions outside of the UK to compete for some customers, especially in Europe and North America, but increasingly other regions too. For example, UK companies with global offices and customers may want to locate some applications in data centres outside of the UK (particularly, the US and EU) to be close to the customer for latency reasons. Geographic reach may be particularly important for ISVs looking to host their services on the infrastructure of cloud providers as they will want to appeal to a broad range of global customers and, for some applications, latency requirements may necessitate locating in regions close to customers. For example, [✂] submitted that it would be difficult to self-supply cloud infrastructure on all of the hyperscalers' geographical regions and it would be difficult to deliver its service from outside

---

[762] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂].
[763] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂].
[764] [✂].
[765] For example, OVHcloud has 37 data centres in 4 continents, Scaleway operates data centres in one region (Paris), and Civo operates data centres in 3 regions. OVHcloud website, OVHcloud datacentres [accessed 19 September 2023]; Scaleway website, Scaleway Datacenter homepage [accessed 19 September 2023]; Civo website, Frankfurt region now live [accessed 19 September 2023].
[766] Oracle response dated 9 December 2023 to the s.174 notice dated 31 October 2022, Part B question 22a.
[767] Analysis of provider websites suggests that AWS, Microsoft, Google and Oracle all offer public cloud services to customers from over 30 geographic regions located around the world. Also, in 2021, IBM and Oracle both operated data centres in 19 countries, AWS in 19 countries, Microsoft in 24 countries, and Google in 21 countries. Synergy Research Group, 2022. 1Q 2022 Hyperscale Market Tracker.
[768] For example, in our qualitative research, one respondent considered that only Google, AWS and Microsoft have the reach to meet the needs of their global business. Context Consulting research report, slide 64.

the hyperscaler regions due to potential latency issues which would be critical for use-cases requiring a [✂].[769] [✂].[770]

6.15 Overall, the hyperscalers are still ahead of Oracle in terms of the total number of data centres available for providing public cloud services, which means that they may be able to offer more availability within geographic regions (e.g. multiple availability zones within regions).[771] Furthermore, the hyperscalers are at an advantage as they have already achieved expansive global networks of larger data centres and have already made the capital investments. They may also be better able to spread the fixed costs of capital investments for further expansion due to the significant scale of their existing cloud customer base.

## Economies of scale in data centres

6.16 In addition to high capital costs, economies of scale can further exacerbate barriers to entry and expansion. The hyperscalers may be able to benefit from economies of scale given the size and global reach of their data centres and larger global customer base.

6.17 Evidence from existing reports suggests that there are significant economies of scale associated with the size of data centres, with costs per kilowatt (kW) of compute capacity decreasing with increases in data centre size – this is the case for all cost categories, especially for energy and operating costs, which together account for 80% or more of annual data centre costs.[772] Additionally, a Microsoft study states that investments in security and reliability also benefit from economies of scale, as these are largely fixed costs.[773] There is also evidence that data centres with higher average rack density experience lower unit costs than data centres with lower rack density.[774]

6.18 On average, the hyperscalers have larger data centres (measured on a MW basis) globally than the small-scale and mid-scale cloud providers, therefore we would expect that they would benefit from economies of scale associated with the size of data centres (see Table 6.1).[775] For example, Microsoft's data centres have an average capacity of [✂] [10-20] MW

---

[769] [✂] response dated [✂] to the s.174 dated [✂], Part A question [✂].
[770] [✂] response dated [✂] to the s.174 notice dated [✂], Part B, question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂].
[771] For example, AWS offers Availability Zones (AZs) which are one or more discrete data centres with redundant power, networking, and connectivity in an AWS Region. AWS explains that "AZs give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center" AWS website. Regions AZ [accessed 19 September 2023].
[772] Energy costs experienced the largest decrease, with a 180 percent difference between energy costs/kW for data centres in the smallest size range compared to the largest. Operating costs showed a 129 percent difference per kW in the largest data centres compared to those in the 500 to 5,000 square foot range. The other cost categories are the physical building and the physical infrastructure. Emerson & Ponemon Institute, 2016, Cost to Support Compute Capacity; results reported in Data Center Dynamics, 2016, Research: Larger data centres make considerable savings on operating costs [accessed 19 September 2023].
[773] Microsoft, 2010, The Economics of the Cloud [accessed 19 September 2023].
[774] Ibid.
[775] The hyperscalers may also experience benefits in terms of customer demand. In the broadcasting space, one ISV [✂] told us that it chooses to run their services on the hyperscalers as it has concerns that smaller cloud providers are unable to meet peaks of processing demand. It noted that smaller providers are good for dedicated workloads but said that it does not want to separate out its product into separate workloads to avoid latency issues. Ofcom / [✂] meeting, [✂].

per data centre globally compared to Oracle's data centres which have an average capacity of [✂] [less than 5] MW.[776] When we compare infrastructure costs to revenue based on information provided by the hyperscalers, [✂].[777] [✂], this trend is consistent with economies of scale, particularly in relation to data centre usage costs.

6.19 A Microsoft study states that economies of scale can also be achieved with data centres in terms of buying power. Operators of larger data centres can get discounts on hardware purchases of up to 30% over smaller buyers.[778] Whilst [✂] told us that it expects the gap in server costs to be between 10-30% when compared to a hyperscaler, [✂] explained that it believes there is enough competition in the server space (Intel, AMD, Nvidia, etc.) to ensure that smaller providers can access components at relatively competitive prices.[779] However, this may not be true during periods of shortages.[780] [✂] told us that in a context of electronic components shortage, suppliers favour clients with larger purchasing volumes (i.e. hyperscalers) [✂].[781]

6.20 As Microsoft acknowledged in response to our interim report, there are also economies of scale associated with pooling of demand to reduce variance and improve the utilisation rate of servers and data centres.[782] There are various sources of demand variability which can lead to under-utilisation – including random variability, time-of-day patterns, industry-specific variability, and variability from different workloads.[783] The larger the pool of customers (both in quantity and variety), the smoother the demand profile which leads to higher utilisation of servers leading to cost efficiencies.[784] Cloud providers with a global network of data centres and a large established customer base may be better able to reduce variability in time-of-day patterns by, e.g. running the same workload for multiple time zones on the same servers. Indeed, on a global scale IBM explained that the need to have a large global infrastructure footprint combined with sufficient scale (in terms of customer base) are the two most important factors influencing utilisation in data centres and profitability.[785]

6.21 Our evidence on differences between the hyperscalers and smaller cloud providers in terms of data centre/server utilisation rates is mixed. [✂][786]

---

[776] Calculated using data on global data centre count and global data centre capacity. Microsoft response dated 9 December 2023 to the s.174 notice dated 21 October 2022, Part B question 26a (Confidential Annex B26); Oracle response dated 9 December 2023 to the s.174 notice dated 31 October 2022, Part B question 22a.
[777] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂].
[778] Microsoft, 2010. The Economics of the Cloud [accessed 19 September 2023].
[779] Ofcom / [✂] meeting, [✂].
[780] It has been reported that the development of the AI industry could lead to a period of chip shortage. See for example CNN Business, 2023. The big bottleneck for AI: a shortage of powerful chips [accessed 14 August 2023].
[781] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[782] Microsoft response to the interim report, paragraph 66.
[783] Microsoft, 2010. The Economics of the Cloud, pages 5-6 [accessed 19 September 2023].
[784] Higher utilisation of servers would lead to cost efficiencies as cloud providers would require fewer servers for the same demand and less power. Microsoft, 2010, The Economics of the Cloud, page 7 [accessed 19 September 2023].
[785] IBM response dated 22 March 2023 to our proposed use of information dated 14 March 2023.
[786] [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], Part B question [✂].

6.22 Both AWS and Microsoft have made public statements highlighting the economies of scale that their cloud businesses enjoy:

a) AWS stated: "…our leadership position helps: scale economies can provide us a relative advantage on capital efficiency."[787]
b) Microsoft stated: "Our cloud business benefits from three economies of scale: datacenters that deploy computational resources at significantly lower cost per unit than smaller ones; datacenters that coordinate and aggregate diverse customer, geographic, and application demand patterns, improving the utilization of computing, storage, and network resources; and multi-tenancy locations that lower application maintenance labor costs."[788]

6.23 Responding to our interim report Microsoft agreed that the upfront investment required to provide IaaS and the substantial economies of scale and scope tend to favour the large-scale players but stated that it would not expect an efficient market structure for IaaS to be fragmented across large numbers of competitors.[789] We acknowledge that scale economies may imply the scope for competition in IaaS is limited to a small number of players.

6.24 Microsoft also noted that, in recent years, cloud adoption has expanded to more specialised business requirements and suggested that this has opened opportunities for existing and new cloud providers to differentiate themselves in different industry and workload verticals, without the need for hyperscaler scale.[790] We recognise that there can be some workload competition (particularly at the PaaS layer) and that market players such as ISVs can start providing cloud services by using the infrastructure of the hyperscalers. However, as discussed later in this section, our evidence indicates that the extent to which this can put competitive pressure on the market leaders may be limited.

6.25 Another hyperscaler respondent [✂] recognised that certain features of its offering involved significant investment, but stated that competitors do not need to replicate its full offering or have the capacity to supply 100% of every customer's requirements in order to compete because many customers multi-cloud.[791] While it may be true that customers use multi-cloud solutions for a limited number of siloed use cases, our evidence suggests that usage of integrated multi-cloud is low and unlikely to increase due to the barriers we have identified in Section 5. Therefore, we disagree that multi-cloud is sufficiently widespread to materially lower barriers to entry and expansion by allowing small providers to compete for a fraction of customer demand and gradually build out. Moreover, as discussed above, some customers want access to a large network of data centres which puts the hyperscalers at an advantage as they have already achieved expansive global networks of larger data centres.

## Innovation in underlying hardware

6.26 Innovation in underlying hardware, such as custom processors and custom hardware accelerators, can drive significant improvements in performance of hardware (e.g. compute, memory and storage) which translates into lower energy use. These efficiencies lead to lower unit costs for cloud providers; as explained above, energy costs account for a significant share of annual data centre costs. For example, AWS's latest generation of Arm-

[787] Amazon, 2014. Letter from CEO to shareholders, page 5 [accessed 28 September 2023].
[788] Microsoft, 2022. Microsoft Annual Report 2022 [accessed 19 September 2023].
[789] Microsoft response to the interim report, paragraph 64.
[790] Microsoft response to the interim report, paragraph 71.
[791] [✂].

based processors (Graviton3) provides up to 25% better compute performance and up to 60% less energy usage.[792] Other examples, of hyperscalers customising hardware are provided in Box 6.2 below.

6.27 These efficiencies can be passed onto customers in terms of lower unit prices and increased performance.[793] This in turn improves the price-performance ratio of the cloud services running on top of it.[794] In other words, innovation in underlying hardware translates into better quality services and value for money for consumers across the entire cloud stack. The market research shows that quality of service and value for money are the two most important factors when choosing a provider.[795] This is also supported by the responses we have received from large enterprises: [✂], [✂], [✂], [✂] and [✂] all considered service quality to be an important factor when choosing cloud provider, with [✂] and [✂] ranking it as their most important consideration.[796] [✂] also considered service quality to be an important factor when choosing cloud provider, alongside price, and said that it focuses on "functionality and performance such as cloud host CPU performance, disc performance, network stability, etc."[797]

6.28 Investments in research and development (R&D) to develop custom hardware involve high fixed and sunk costs and technical expertise, which may act as a barrier to entry and expansion for smaller cloud providers. The more customers a company has, the more it can spread these fixed costs. Furthermore, the hyperscalers use these innovations internally and often first customise hardware for their non-cloud businesses.[798] For example, Google's TPU chips were initially used exclusively for its non-cloud services, e.g. Google Search.

6.29 Access to similar innovations from third-party suppliers could reduce barriers to competing with cloud providers who have invested in optimising their underlying hardware. For example, we see evidence of Microsoft and Oracle accessing innovations from chip makers

---

[792] Compute performance is measured based on a comparison to AWS Graviton2 processors and energy savings are based on "Graviton-based instances using up to 60% less energy for the same performance than comparable EC2 instances". We understand "comparable EC2 instances" to mean EC2 instances run using other processors that also available on EC2. AWS website, AWS Graviton Processor [accessed 19 September 2023].

[793] For example, according to AWS's website AWS Graviton based instances deliver up to 40% better price performance over comparable current generation x86-based instances for a broad spectrum of workload. Price performance is calculated based on: "20% lower cost and up to 40% higher performance for M6g, C6g, and R6g instances over M5, C5, and R5 instances respectively, based on internal testing of workloads with varying characteristics of compute and memory requirements." Our understanding is that costs are estimated using price per hour for Graviton versus current generation x86-based instances. AWS website, AWS Graviton Processor [accessed 19 September 2023].

[794] This process may be delayed for third-party services as an ISV or third-party cloud provider may not be able to fully test and update their applications until sometime after a new API release from the underlying infrastructure provider. We discuss this in more detail in Section 7.

[795] The market research found that customers care most about quality of service (most popular reason) and value for money (second most popular reason). Context Consulting research report, slide 68.

[796] [✂] response dated [✂] to our customer questionnaire; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our customer questionnaire; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

[797] [✂] response dated [✂] to our customer questionnaire.

[798] For example, Google explains that "the state-of-the-art capabilities you see in our products such as Search and YouTube are made possible by Tensor Processing Units (TPUs), our custom machine learning (ML) accelerators." Google Cloud blog, 2022, Google Cloud unveils world's largest publicly available ML hub with Cloud TPU v4, 90% carbon-free energy [accessed 19 September 2023].

such as Nvidia and Ampere to compete directly with other providers, including the hyperscalers:

a) In 2022, Azure launched a preview of virtual machines on its cloud service powered by Arm-based server CPUs from startup Ampere which "compete directly with [AWS's] Graviton".[799]
b) Oracle has also been using Ampere CPUs and powerful graphics processors and GPUs from chipmaker Nvidia.[800] There are examples of Oracle using Ampere chips to improve its services and offer competitive prices.[801]

6.30 Overall, the ability for cloud providers to externally source innovative hardware may lower the barriers to entry and expansion from the need to invest in solutions that can increase the efficiency of the infrastructure required to provide cloud infrastructure services.[802] However, providers who are unable to invest in such innovations may still experience some cost disadvantages due to the need to purchase such solutions from external providers.

**Box 6.2: Examples of the hyperscalers customising hardware for data centres**

AWS has customised hardware for its data centres, including i) Arm-based processors (Graviton CPUs),[803] ii) Nitro smart NICs/data processing units,[804] and iii) Inferentia[805] and Trainium[806] ML accelerators. [✂].[807] AWS launched its first Arm-based processor (Graviton CPU) in November 2018[808] and has since launched Graviton 2 and Graviton 3.[809] AWS's development of Graviton CPUs alone has resulted in substantial price performance benefits for its customers relative to x86-based processors (provided by Intel and AMD).[810]

---

[799] Electronic Design, 2022, Microsoft Taps Ampere's Arm CPUs for New Cloud Service [accessed 19 September 2023].
[800] "Oracle Cloud Infrastructure (OCI) will expand its offering of online accessible computers running NVIDIA's powerful A100 graphics processors, connected by fast networking and aimed at industries including banking, healthcare, and manufacturing. Oracle also plans to offer the chipmaker's upcoming H100 "Hopper" GPUs, which can shrink AI model training time from 7 days to 20 hours for some workloads." Oracle Connect, 2022, Oracle Cloud adds NVIDIA chips, software to speed enterprise AI uptake [accessed 19 September 2023]. Furthermore, Oracle has made significant investments in Ampere. Protocol, 2022, Oracle has pumped more than $400 million into chip startup Ampere [accessed 19 September 2023].
[801] Oracle website, Oracle Unlocks Power of Arm-based Processors at One Cent per Core Hour, Expanding Ecosystem, and Speeding App Development [accessed 19 September 2023].
[802] A hyperscaler [✂] agreed with this finding. [✂].
[803] AWS website, AWS Graviton Processor [accessed 19 September 2023].
[804] AWS website, AWS Nitro System [accessed 19 September 2023].
[805] AWS website, AWS Inferentia [accessed 19 September 2023].
[806] AWS website, AWS Trainium [accessed 19 September 2023].
[807] [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂].
[808] Graviton CPUs are built around ARM cores and making extensive use of customised silicon. AWS News Blog, 2018, New – EC2 Instances (A1) Powered by Arm-Based AWS Graviton Processors [accessed 19 September 2023]
[809] AWS website, AWS Graviton Processor [accessed 19 September 2023].
[810] In 2020 AWS released Graviton2 which generates 20% lower cost and up to 40% higher performance over comparable current generation x86-based instances for a broad spectrum of workloads. In 2021, AWS announced Graviton3 which offers even better performance than Graviton2 for additional computing workloads, including three times better performance compared to AWS Graviton2 processors for machine learning. AWS website, AWS Graviton Processor [accessed 19 September 2023]; also see AnandTech, 2020, Amazon makes Graviton2 AWS instances available [accessed 19 September 2023].

It has been reported that Microsoft and Google will similarly customise their own Arm-based processors for their respective cloud services.[811] [✂].[812] [✂].[813]

Google has already customised hardware accelerators, Tensor Processing Units, (TPU) which are designed to speed up ML and has made these available to customers on Google Cloud since 2018 (Cloud TPU).[814] Google initially developed its TPU for its data centres hosting Google Search, Street View, Google Photos and Google Translate and has been using them internally since 2015.[815] Its first TPU delivered 15-30 times higher performance and 30-80 times higher performance-per-watt than contemporary CPUs and GPUs. Google explained that these advantages helped many of Google's (non-cloud) services run state-of-the-art neural networks at scale and at an affordable cost.[816] Google explains that its cloud customers "can tap into the same custom-designed machine learning ASICs (application-specific integrated circuits) that power Google's Search, YouTube, and LaMDA AI model" to speed up their own machine learning models.[817]

## Benefits from sharing cloud infrastructure with non-cloud businesses

6.31 AWS, Microsoft and Google's cloud infrastructure businesses are likely to benefit from being part of large tech conglomerates with significant digital non-cloud businesses.

6.32 They can use their own public cloud services internally in other parts of their non-cloud businesses. In doing so, the non-cloud businesses can act as large 'anchor tenants', which can guarantee a minimum level of demand for their cloud services. This could make it easier to realise economies of scale, and increase the expected return on investments. Examples of the hyperscalers supplying cloud to their non-cloud businesses include:

a) Amazon's Consumer Business (including Amazon Prime, Amazon Prime video etc.) migrated 75 petabytes of internal data stored in nearly 7,500 Oracle databases to multiple AWS database services in 2019.[818]
b) Microsoft uses Azure to power its Bing search engine, Xbox Live services and has migrated most of its Office365 services to Azure.[819]

---

[811] Bloomberg, 2020, Microsoft Designing Its Own Chips for Servers, Surface PCs [accessed 19 September 2023]; and The Register, 2023, Taking notes from AWS, Google prepares custom Arm server chips of its own [accessed 19 September 2023].
[812] [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂].
[813] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[814] Google Cloud website, Cloud TPU [accessed 7 March 2023].
[815] Google Cloud blog, 2017, An in-depth look at Google's first Tensor Processing Unit (TPU) [accessed 19 September 2023].
[816] Ibid.
[817] Ibid.
[818] The migration involved 100 teams in Amazon's consumer facing business including Amazon Prime, Amazon Prime Video etc. as well as internal teams. Oracle databases were replaced with several AWS databases including Amazon DynamoDB, Amazon Aurora, Amazon Relational Database Service (RDS), and Amazon Redshift. Amazon blog, 2019, Migration Complete – Amazon's Consumer Business Just Turned off its Final Oracle Database [accessed 19 September 2023].
[819] ZDNET, 2021, Microsoft moves closer to running all of its own services on Azure [accessed 19 September 2023].

c) Google announced plans to move parts of YouTube to its Google Cloud platform in 2021.[820]

6.33 Whilst some of the hyperscalers' major non-cloud services (e.g. Google Search) may currently be hosted separately to its public cloud services, it can still use these services to realise economies of scale in the underlying infrastructure (e.g. data centres, hardware and networks can be shared across their cloud and non-cloud businesses). This means that the hyperscalers can invest in larger data centres and realise economies of scale. Furthermore, investments in innovation in both the underlying infrastructure and cloud services benefit both cloud and non-cloud businesses. This can create economies of scope as skilled technical resources, and the fixed costs of R&D are spread across a range of cloud and non-cloud products and services. For example, Google's TPU chips were initially used exclusively for its non-cloud services including Google Search which it later made available to customers on its public cloud.[821]

## Conclusion

6.34 In our interim report we said that the need to invest in physical infrastructure is likely to create significant barriers to entry and expansion for cloud providers. There is some scope to phase investments and lease data centres, for example by entering with a single leased data centre in one region and gradually increasing data centres in number, size and geographical reach. We observe that some small-scale cloud providers have entered the market and are increasing their physical infrastructure. However, they have significantly fewer data centres and in fewer regions. Our views in this respect therefore remain unchanged.

6.35 We still consider that a global network of data centres is required to compete effectively for some customers, which takes time and significant capital expenditure to achieve. As such, the hyperscalers are ahead of other cloud providers having already made the capital investments and established expansive global networks of large data centres. Of the mid-scale cloud providers, Oracle has been able to make the most progress in this area. Oracle has been able to establish a larger network of data centres having only entered in 2016, although it remains significantly behind the hyperscalers in terms of the size and density (i.e. number) of data centres.

6.36 Furthermore, the hyperscalers are likely to have an advantage due to the significant scale of their existing cloud customer base (including an anchor tenant associated with their non-cloud business). Relative to other cloud providers, they are likely to benefit from economies of scale given the size of their data centres and the ability to achieve higher utilisation rates in their data centres.[822] The hyperscalers may also be better able to spread the fixed costs of their capital investments for further expansion.

6.37 Investment in physical infrastructure is not a barrier to entry and expansion for the supply of PaaS, as ISVs can typically use the infrastructure (IaaS services) of other cloud providers to offer their services rather than using their own infrastructure.

---

[820] DCD, 2021, Google to migrate parts of YouTube to Google Cloud [accessed 19 September 2023].
[821] Google Cloud blog, 2022. Google Cloud unveils world's largest publicly available ML hub with Cloud TPU v4, 90% carbon-free energy [accessed 29 September 2023].
[822] We do not have precise estimates of the global customer base for hyperscalers and other cloud providers, but in Annex 2 we discuss the annual global revenue of major cloud providers and show that the hyperscalers generate higher public cloud revenue than other cloud providers. This indicates that hyperscalers' global public cloud customer base is very likely to be larger than that of other cloud providers.

# Investment in product portfolios

6.38 In this subsection, we examine the extent to which the need to invest in a wide breadth of products may act as a barrier to entry and expansion. The subsection is structured as follows:

a) First, we discuss the importance of breadth and quality of offering and summarise the offerings of the main cloud providers.
b) Second, we examine whether small-scale and mid-scale cloud providers can match the rate of innovation and range of products and features of the hyperscalers.
c) Third, we examine whether network effects may act as a further barrier to entry and expansion.

## The breadth and quality of providers' product ranges is an important parameter of competition

6.39 As detailed in Section 4, being able to offer a broad range of high-quality services is important for cloud providers to be able to attract customers and differentiate themselves.

6.40 Breadth of product range and the quality of those services is important for many customers because within their organisations they are likely to have a variety of use-cases that require different types of products. This is consistent with the market research, which found that the top five reasons selected by customers for choosing a particular provider included service quality (top reason) and number of features (ranked fifth).[823] The market research found that "service quality" and "number of features" were cited as important factors when choosing a provider in 39% and 31% of cases respectively.[824] Furthermore, 84% of users have four use cases or more.[825] In line with this, many of the customers responding to our customer questionnaire (e.g. [✂]) noted they value range of services when picking a cloud provider.[826] This is further reflected in customer purchase data, with the average hyperscaler customer purchasing multiple services.[827]

6.41 Product range is also important because a cloud provider with a broader product range will be able to attract a broader range of customers, and therefore ultimately attract a greater number of customers. Certain customers have specialist (e.g. industry-specific) use-cases where they may benefit from providers having a wide range of products that can be combined into tailored solutions, or cloud products that are specialised for their needs. For example, a customer ([✂]), told us that it currently procures cloud services from AWS and that the top three most important factors it considers when procuring cloud services is

[823] Context Consulting research report, slide 69; and Context Consulting research data tables, Q25.
[824] Context Consulting research report, slide 69; and Context Consulting research data tables, Q25.
[825] Context Consulting research report, slide 35.
[826] [✂] response dated [✂] to our customer questionnaire. Ofcom / [✂] meeting, [✂]. [✂] response dated [✂] to our customer questionnaire. [✂] response dated [✂] to our customer questionnaire. [✂] response dated [✂] to our customer questionnaire.
[827] [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂])]; [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], Part B question [✂].

service quality, reputation (including broad capabilities of AWS) and access to software only available on a specific cloud provider's platform.[828]

6.42 Customers also value their providers offering them access to the latest innovations – for example, the hyperscalers have told us that 'rate of innovation' is an important capability customers consider when choosing cloud providers.[829] This is also consistent with the market research which found that service quality and number of features are important factors for consumer choice of cloud provider.[830] It is therefore necessary for cloud providers to maintain a high rate of innovation to develop the range and quality of their products.

6.43 This explains why the hyperscalers have developed ecosystems of large product portfolios that span the full cloud stack, which allows customers to source products from one place and to easily combine them to build their IT solutions.

6.44 In this regard, AWS markets itself as "the most comprehensive" cloud platform, by offering "over 200 fully featured services" which it states is "more services, and more features within those services, than any other cloud provider".[831] Data we have gathered from providers (shown in Table 6.3 below) indicates that AWS does indeed offer the most services, closely followed by Microsoft and Google. Oracle and IBM also appear to offer a considerable number of products, although their product-counts are lower than the hyperscalers. In terms of breadth of offering, Figure 6.4 below indicates that AWS, Microsoft and Google are offering products across similar numbers of product categories.

**Table 6.3: Number of cloud infrastructure products by provider – hyperscalers and mid-scale providers**

| | Number of products |
|---|---|
| **AWS** | 220+ |
| **Microsoft** | 200+ |
| **Google** | 190+ |
| **Oracle** | 110+ |
| **IBM** | 90+ |

*Source: Ofcom analysis of IaaS and PaaS products listed on provider websites (September 2023).*

[828] [✂] response dated [✂] to our customer questionnaire. [✂] outlined that it considers technical capabilities, which include the range of services offered and their functionality (including the geographical availability of such services). [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[829] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[830] Context Consulting research report, slide 69; and Context Consulting research data tables, Q25. The list of potential response options didn't include "innovativeness" as a response option. Therefore, we would not expect "innovativeness" to have been specifically mentioned, but instead is likely to feature in consumers perceptions of factors listed such as quality of service and number of features. The top five reasons (in order) all arguably include elements of innovation: service quality; best value for money; supplier reputation; proposed level of security; and number of features.
[831] AWS website, What is AWS [accessed 19 September 2023].

**Figure 6.4: Number of product categories served, based on provider websites**[832]

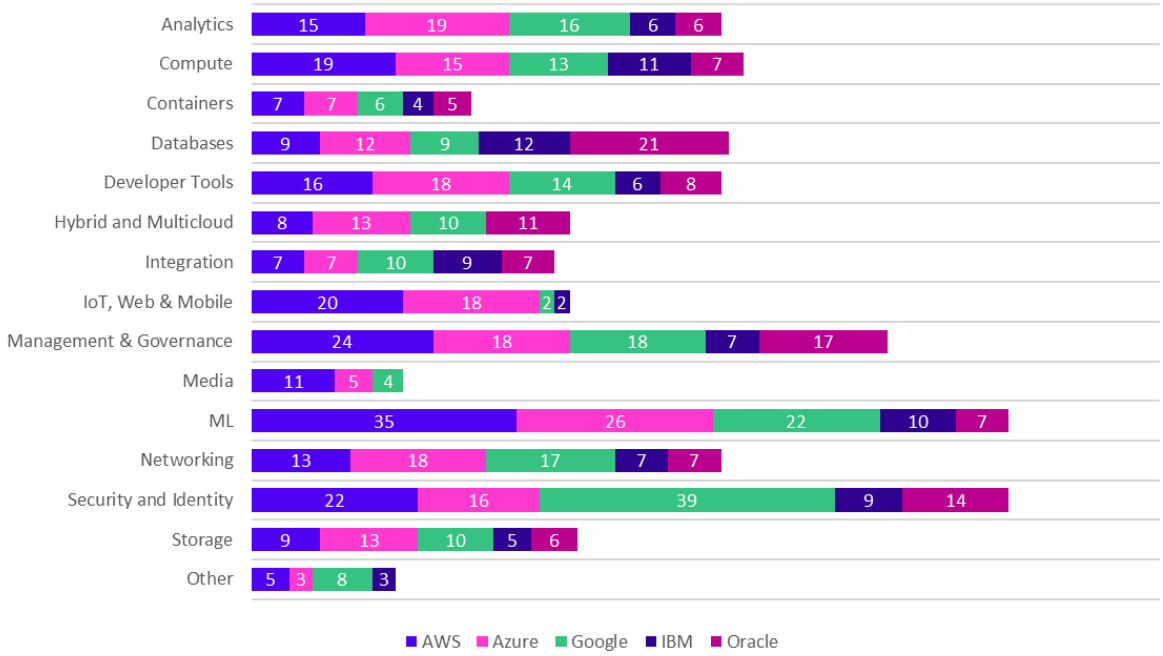


*Source: Ofcom analysis of provider websites and mapping of providers' services to IaaS and PaaS categories (September 2023). Note: IBM does not have a dedicated hybrid/multi-cloud category. Instead, related services are captured in the Integration, Networking and Containers categories.*

6.45 One respondent to our interim report ([✂]) stated that providers can and do start with a narrower range and expand as their business grows, pointing to the mid-scale providers Oracle and IBM as having achieved relatively broad product ranges this way.[833] Indeed, Oracle's efforts to broaden its range of cloud services and features have been acknowledged by industry experts ([✂]) as bringing it closer to the market leaders in terms of hyperscale cloud capabilities.[834] However, as observed in our interim report and as illustrated by Figure 6.4 above, Oracle still offers a slightly narrower range and fewer products within most categories.[835] For example, Oracle does not have PaaS products in the media category. Also, although Oracle has a presence in the machine learning and AI ("ML") category, its services within this category are more general purpose (e.g. Oracle offers a generic data anomaly detection service similar to AWS, but AWS also has bespoke services for specific industries, e.g. healthcare and industrial applications).[836] On the other hand, we see that Oracle offers more products in the databases and hybrid and multi-cloud categories, which likely reflects its areas of strength from its offering in on-premises business software.

[832] AWS website, AWS; Microsoft website, Microsoft; Google website, Google; Oracle website, Oracle; IBM website, IBM.
[833] [✂].
[834] [✂].
[835] These are based on how many categories of products each of these providers claims to offer on their websites. Whilst these numbers will depend on how each provider categorise their products, we found that product categories are broadly consistent across provider, and therefore consider that combined with product count, they provide an indication of breadth of product range.
[836] Oracle website, AI services [accessed 19 September 2023]; and AWS website, AI services [accessed 19 September 2023].

6.46 IBM offers a narrower range of services across IaaS and PaaS and fewer products than the hyperscalers in most categories. It told us that its current cloud strategy is to offer a hybrid cloud platform which allows their clients to span workloads across many cloud services providers, including all of the hyperscalers, and focus on complex mid- and back-office workloads where it brings differentiated value.[837]

6.47 In contrast, we observe that small-scale cloud providers do not have the same breadth of product range, as indicated by the lower numbers of products offered by the providers we have looked at in Table 6.5 below. In addition, it is our understanding that the product ranges of these small-scale cloud providers are focussed on IaaS services which are generally more commoditised.

**Table 6.5: Number of cloud infrastructure products by provider – selected small-scale providers**

| | Number of products |
|---|---|
| **OVHcloud** | 17 |
| **Digital Ocean** | 23 |
| **Scaleway** | 25 |

*Ofcom analysis of IaaS and PaaS products listed on provider websites*

6.48 We observe that the hyperscalers and Oracle[838] are consistently adding new products and features across the entire cloud stack to expand their broad portfolios of services. This requires continuous investment and innovation:

a) Our analysis of data collected from the hyperscalers in response to our statutory information requests indicates that they each added [✂] new services per annum in net terms to their product portfolios (i.e. factoring in service discontinuations) between 2020 and 2022.[839] For example, AWS released between [✂] new services in each year between 2020-22.[840] As well as new services, the hyperscalers also appear to be releasing many new features within services on a regular basis. For example, "The Stack counted over 119 new AWS services and features landing during the cloud hyperscaler's re:Invent 2022 conference."[841]

b) [✂] told us that the hyperscalers maintain a "very high rate of innovation", providing an example of when a customer using AWS IaaS decided to use CloudFoundry (an open-

[837] IBM response dated 22 March 2023 to our proposed use of information dated 14 March 2023.
[838] [✂] explains that Oracle continues an impressive year-over-year pace of feature velocity that brings it closer to the market leaders in terms of hyperscale cloud capabilities. [✂] explains that if the pace continues, Oracle will meet or exceed some of the providers in the Leaders quadrant in terms of capabilities within the foreseeable future. [✂].
[839] The equivalent figures in gross terms, i.e. excluding service discontinuations, are slightly higher. In any given year, the hyperscalers' new public cloud services typically represent [✂]% of their total service count, but this proportion can sometimes be higher. Ofcom analysis of data collected from the hyperscalers in response to s.174 notices.
[840] Count of new public cloud services (IaaS and PaaS) only by year (net). Ofcom analysis of data collected from the hyperscalers in response to s.174 notices.
[841] The Stack, 2022. 119 new AWS services and features in 30 words each [accessed 19 September 2023].

source PaaS) to develop new features initially not available from AWS, but were made available by the time the customer developed the features.[842]

6.49 Based on the evidence above, we consider the need for providers to maintain and invest in the breadth and quality of their product portfolios to be a barrier to entry and expansion. While we recognise that it is possible for providers to build product range gradually, the product ranges of small-scale and mid-scale cloud providers remain narrower than those of the hyperscalers – even for Oracle, who is continuously investing to add new products and features. Given that many customers value access to broad product portfolios, they are more likely to choose a hyperscaler than any of the smaller cloud providers. Although in principle it is possible for a customer to access range by combining the offerings of different providers, this is limited due to the barriers to multi-cloud identified in Section 5.[843] As such, barriers to multi-cloud can increase barriers to entry and expansion, as smaller cloud providers are less able to grow scale by competing for certain components of customer demand – and this is reflected in market shares.

6.50 Additionally, as well as investing in building a product range, it may be necessary for providers to invest in increasing or maintaining awareness among customers of their capabilities. The evidence we gathered indicates that some customers view cloud providers other than AWS and Microsoft as offering less comprehensive offerings. For example:

a) Google: Our qualitative research indicated that some customers did not consider choosing Google as their cloud provider because they view Google as serving more niche use-cases or customer types. One respondent said: "With Google, we rightly or wrongly have a view of it mainly being for Big Data and for maybe co-development related to having all your data in one place and being smart about different tools that can analyse different parts of that".[844]
b) Oracle: ([✂]) has noted that there is still a perception amongst customers that Oracle is not a general-purpose cloud provider and is not positioned for adoption by midmarket enterprises and small and medium-sized businesses, suggesting that it may take some time and effort for Oracle to reposition itself.[845]

6.51 As outlined above, in response to our interim report one hyperscaler ([✂]) noted that for PaaS and similar cloud services, scale is less relevant because competition centres on catering to more specialised business needs. As such, it suggested that smaller players can compete on an equal footing with larger players.[846] In the interim report we acknowledged that, in contrast to cloud providers, many ISVs have entered the market by offering PaaS services only and with products that focus on a single PaaS product category. While there are many more ISVs than cloud providers, our analysis suggests no single ISV has a share in PaaS greater than 5%.[847] ISVs tend to specialise in providing one type of PaaS category (e.g. data management services). For example, in 2022, most ISVs tracked by IDC were only active

[842] Ofcom / [✂] meeting, [✂].
[843] One hyperscaler ([✂]) said it believes that increasing multi-cloud and ensuring interoperability would mitigate against the perceived need for each cloud provider to maintain a broad range of products. It suggests that product range and innovation do not amount to material barriers to entry and expansion, as long as the industry is willing to commit to open standards and a high degree of interoperability. [✂].
[844] Context Consulting research report, slides 62-63.
[845] [✂].
[846] [✂] confidential response to the interim report, paragraph [✂].
[847] Ofcom analysis of data provided in response to our information requests and data from Synergy and IDC.

in one or two of IDC's seven PaaS product categories.[848] This contrasts with the hyperscalers, which in 2022 were present across all seven of IDC's PaaS categories.

6.52 While this supports the view that it is possible to provide PaaS products without having a broad portfolio of products, this does not mean that ISVs are able to compete on an equal footing with cloud providers. Individual ISVs cannot compete with cloud providers for solutions that combine products across several PaaS and/or IaaS categories. There are also some technical and commercial factors that may disadvantage ISVs when competing with first-party products (discussed in Sections 5 and 7). As a result, many customers favour taking-up first-party products from their primary cloud provider, and ISVs cannot constrain the IaaS elements within customers' workloads.

## Providers broaden their range of products and features through R&D and acquisitions

6.53 Cloud providers can develop their product range and maintain a high rate of innovation through internal R&D and acquisitions. This raises barriers to entry and expansion in several ways:

a) There are high fixed and sunk costs and economies of scope associated with internal R&D.
b) Larger cloud providers may have an advantage due to the skills and experience acquired by their staff (experience curve effect) and find it easier to attract scarce technical skills.
c) Larger cloud providers may have greater financial capabilities to acquire firms and/or make significant equity investments in firms that are innovating and/or have access to specific data and expertise.

### Investment in R&D

6.54 Innovating and developing new products and features internally requires significant investments in R&D involving high fixed and sunk costs. Our financial evidence on R&D spend suggests that the investment required to compete at scale and pace with the hyperscalers may be out of reach for smaller cloud providers. [✂].[849]

6.55 Furthermore, the development of innovative software services (e.g. in PaaS) has high fixed R&D costs but low marginal costs. The more customers a company has the more it can spread these fixed costs.[850, 851]

[848] By PaaS product category here we are referring to IDC's seven 'secondary markets' for PaaS which are analytics and business intelligence, AI platforms, data management, integration and orchestration, application development, software quality and life cycle, and application platforms. IDC, Public Cloud Services Tracker 2022 H2 (published April 2023).
[849] Based on analysis of data collected via responses to s.174 notices.
[850] Traditional endogenous growth models argue that due to the nonrival yet partially excludable nature of ideas or designs, R&D efforts are subject to economies of scale – once the initial innovation investment is made, the resulting new design or idea can be used as often and as widely as desired. See Romer, M., 'Endogenous Technological Change', *Journal of Political* Economy, 1990, pages 74-78
[851] Barwise and Watkins (2018) argue that the same premise can be applied to software and digital content, which have high fixed development costs and low-to-zero copying and distribution costs. Barwise, P. & Watkins, L. '*The evolution of digital dominance: how and why we got to GAFA. In: Digital Dominance: The Power of Google, Amazon, Facebook, and Apple*.' Oxford University Press, New York, 2018, pages 21-49.

6.56 The hyperscalers also benefit from economies of scale and scope associated with being part of a large tech conglomerate, where the fixed costs of R&D can be spread across a range of cloud and non-cloud products and services (as explained earlier in this section).

6.57 One hyperscaler ([✂]) said it recognises that internal R&D entails a degree of fixed and sunk costs, but this is a common feature across many "innovation driven industries", including the wider tech sector, which it said has one of the most dynamic financing incubation infrastructures. It [✂] said that smaller cloud providers have access to significant amounts of financing and tech expertise.[852] However, in relation to investment more broadly, another respondent [✂] argued that it is difficult to attract investors if a provider is perceived to be competing against AWS and Microsoft.[853]

## Access to skilled labour

6.58 Developing new products and features internally requires access to teams of technical staff with the appropriate knowledge and highly specialised expertise. As a result, cloud providers that have entered the market earlier may have a competitive advantage due to the experience curve effect (i.e. its staff have already developed the internal know-how and technological expertise). This may allow them to be more efficient by developing services faster and across a wider range of products and services. Our evidence [✂].

6.59 The hyperscalers may have a further advantage due to their major non-cloud digital businesses, which may contribute to their access to highly skilled technical teams and engineers which they may be able to utilise across cloud and non-cloud businesses or re-purpose to cloud. They may also be better able to attract highly technical staff given their brand reputation in non-cloud markets. Microsoft, Oracle and IBM's existing on-premises software businesses may also mean that they have easier access to skilled technical staff for the creation of software services for the cloud.

6.60 In the interim report we noted that smaller providers and entrants may find it difficult to attract highly specialised technical staff when trying to build their technical expertise, which we said was a scarce resource. One cloud provider ([✂]) disagreed with our characterisation of technical skills and experience as scarce and stated that the tech industry enjoys one of the largest pools of highly skilled talent. It also pointed to layoffs in the tech industry at the back end of 2022 and throughout 2023 as evidence that access to skilled labour is not a significant barrier to entry and expansion.[854] We recognise layoffs in the tech sector could have created some slack in the labour market, though there is evidence to suggest the impact on cloud was minimal.[855] Regardless of labour market conditions, the hyperscalers may be more able to attract highly specialised technical profiles, given their financial capabilities to offer competitive salaries and established reputations as leaders in their field. Indeed, a smaller cloud provider ([✂]) explained that a new entrant might find it difficult to attract highly specialised technical staff where there are staff shortages in the cloud market, and given that the hyperscalers offer extremely comfortable hiring conditions that new entrants might not be able to provide.[856]

---

[852] [✂].
[853] [✂] 'Name Withheld 1' response to the interim report, Q5.2.
[854] [✂].
[855] See, for example CRN, April 2023, AWS Confirms Layoffs Impacting 'Single Digit Percentage' Of Employees [accessed 13 September 2023]
[856] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

6.61 We note that is possible for cloud providers to secure talent by acquiring other companies.[857] This may lower the barriers faced by smaller providers to attract highly skilled specialised staff, though this approach may be inaccessible for the smallest providers if they lack the financial capabilities to do so.

## Acquisitions

6.62 Investments in innovative services can either be developed internally or achieved through acquisitions of external companies that have successfully innovated. The data we have collected suggests that acquisitions are an important source to build product range in cloud services and therefore an important source of growth for cloud providers in the cloud infrastructure services.[858]

6.63 We collated cloud-related acquisition data from five of the largest cloud providers (AWS, Microsoft, Google, IBM, and Oracle) for the period 2018 – 2021. The number of acquisitions made in this period ranges from 11 ([✂]) to 23 ([✂]). [✂] and [✂] have made fewer acquisitions than [✂] but more than [✂] and [✂]. [✂]. As mentioned in Sections 3 and 4, AWS has also made important acquisitions in the past, such as its acquisition of Elemental Technologies in 2015, which increased its media capabilities.

6.64 We note that despite their much lower revenue shares in cloud, IBM and Oracle are large, well-resourced technology companies. The spend of these companies and the hyperscalers over this period on individual acquisitions ranges from a few tens of millions to tens of billions of pounds, and small-scale companies that compete in cloud may not be able to match this level of investment.

6.65 We also examined some recent case studies which illustrated how these providers make use of acquisitions to innovate and release new products:

a) Google is investing heavily in cybersecurity and has acquired a number of companies in this area, most notably Mandiant for $5.4bn in 2022,[859] which has enabled it to expand its Chronicle Security Operations software suite.[860] [✂].[861] This is reflected in its $2.6bn acquisition of data analytics start-up Looker in 2020.[862]
b) Oracle recently acquired the US electronic health records company Cerner in 2022 for $28bn and has integrated it into its new healthcare business unit with plans to continue modernising solutions in the healthcare space.[863]
c) Microsoft acquired Nuance, a provider of speech recognition and AI solutions, in 2022 for $20bn.[864] In a press release, Microsoft said that the acquisition represented the

[857] For example, in 2015 Amazon acquired chip designer Annapurna Labs for an estimated $350-400m. The "Annapurna Labs team" is responsible for building innovation in silicon and software for AWS customers. See DataCenter Knowledge, January 2015, Amazon Buys Stealthy Israeli Chip Startup Annapurna Labs; and Amazon website, Annapurna Labs [accessed 19 September 2023].
[858] Based on analysis of data collected via responses to s.174 notices.
[859] TechCrunch, 2022, Google closes $5.4B Mandiant acquisition [accessed 19 September 2023].
[860] TechCrunch, 2022, Google looks to boost its security cred in the cloud [accessed 19 September 2023].
[861] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[862] TechCrunch, 2020, Google closes $2.6B Looker acquisition [accessed 19 September 2023].
[863] Financier Worldwide, 2022, Oracle agrees to acquire Cerner for $28bn [accessed 19 September 2023].
[864] TechCrunch, 2022, After clearing all regulatory hurdles, Microsoft closes $20B Nuance deal [accessed 19 September 2023].

latest step in its industry-specific cloud strategy, and that it plans to augment its own healthcare offering (Microsoft Cloud for Healthcare) with Nuance solutions.[865]

d) Microsoft notes that as the mobile network industry moves to 5G operators are increasing their uptake of cloud and has made strategic acquisitions (Affirmed Networks and Metaswitch Networks) in this area.[866]

e) IBM's acquisition of open-source software company Red Hat in 2019 for $34bn has enabled it to establish a strong presence in the hybrid multi-cloud space.[867] [✂].[868]

6.66 These examples illustrate how the hyperscalers and mid-scale providers (i.e. Oracle and IBM) have been able to use acquisitions to expand the range of services they can offer and broaden the industries they can appeal to.

### Benefits from access to large volumes of data in non-cloud businesses

6.67 AWS, Microsoft and Google may also have some benefits from being part of large tech conglomerates with significant digital non-cloud businesses. They may be able to use data gathered from other segments of their business (including customer data), to increase the quality and range of their cloud services and potentially better target their cloud services to customers.

6.68 For example, Google uses data from Google Maps, to provide its cloud based 'Google Maps platform' which is now tightly integrated with Google Cloud services and tools.[869] [✂].[870] [✂] suggest that the "hyperscalers are able to leverage their position from the B2C segment into B2B, as they have access to a significant volumes of data (including consumer data) which may often give them a competitive advantage in the provision of B2B cloud solutions".[871]

## Network effects

6.69 Cloud providers can add breadth and variety to their product range by attracting ISVs to offer services on top of their cloud infrastructure. This benefits customers by allowing them to select ISVs' services and combine them with the cloud providers' first-party services when building their cloud solutions.

6.70 Indirect network effects may exist for ISVs' services as the benefit to customers of using a certain cloud provider may increase with the volume and quality of ISV services they can access on that cloud. Similarly, the benefit to ISVs of making their services compatible with

---

[865] Microsoft press release, 2021, Microsoft accelerates industry cloud strategy for healthcare with the acquisition of Nuance [accessed 19 September 2023].

[866] Microsoft website, 2020, Microsoft announces definitive agreement to acquire Metaswitch Networks, expanding approach to empower operators and partner with network equipment providers to deliver on promise of 5G [accessed 19 September 2023]; Microsoft website, 2020, Microsoft announces agreement to acquire Affirmed Networks to deliver new opportunities for a global 5G ecosystem [accessed 19 September 2023].

[867] Red Hat website, 2019, IBM closes landmark acquisition of Red Hat for $34 billion; defines open, hybrid cloud future [accessed 19 September 2023]; Softchoice, 2019, 4 takeaways from IBM's Red Hat Acquisition [accessed 19 September 2023].

[868] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

[869] Google website, Google Maps Platform now integrated with the GCP Console [accessed 19 September 2023].

[870] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

[871] [✂] response dated [✂] to the s.174 notice dated [✂], page [✂].

the cloud of a particular provider may increase with the number of users they can access with that provider.

6.71 The presence of indirect network effects can act as a barrier to entry and expansion for smaller cloud providers. This is because it is costly to develop PaaS solutions for different clouds, such that smaller cloud providers may find it more difficult to attract ISVs onto their clouds due to their small user base. This in turn may make it more difficult for them to attract customers, creating a vicious cycle.

6.72 In addition to this, the existence of indirect network effects may act as a mechanism which favours hyperscalers with a first mover advantage in some or all segments of cloud infrastructure. This is because it takes time to enable PaaS software on different clouds due to their technical differentiation. As a result, ISVs may design their services to be compatible with one hyperscaler at a time based on their popularity, which may reinforce the market position of hyperscalers that are more popular in some or all segments of the market.

6.73 In its response to our CFI, Microsoft noted that, unlike other IT settings,[872] indirect network effects are largely absent in cloud as customers can pick and choose solutions across clouds to build compelling applications which means they do not necessarily care about range of services offered within a specific cloud.[873] However, based on the evidence we have received for our interim report and for this report, we consider that network effects are in fact an important feature of the cloud market. In particular:

a) As noted above, the breadth of product range and the quality of those services is important for many customers, because within their organisations they are likely to have a variety of use-cases that require different types of products.[874]
b) Evidence from the hyperscalers and smaller cloud providers indicates that cloud providers compete to attract ISVs to their clouds to meet users' demand for range of services. For example, Microsoft submitted that cloud providers compete by making available the broadest and most powerful set of functionality possible for developers to create their own applications and services for both internal and external use.[875] Similarly, OVHcloud said that it is in its commercial interest to support the development of these services that are interoperable with its own since such complementarity is highly valued by its customers.[876] Consistent with this, as discussed in Section 4, the hyperscalers offer commercial mechanisms (co-selling) to aid visibility and distribution of ISV services.
c) The evidence from ISVs indicates that they take into account the size of customer base when choosing a provider. For example, three popular ISVs ([✂]) said that expected

[872] 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.
[873] Microsoft response to the CFI, page 12.
[874] This is consistent with evidence from [✂], who submitted that, while less of a barrier compared to cloud credits and egress fees, the larger range of services offered by the hyperscalers demonstrates their commercial advantage compared to smaller cloud providers. Ofcom / [✂] meeting, [✂].
[875] Microsoft response dated 18 November 2022 to the s.174 dated 21 October 2022, Part A question 31.
[876] OVHcloud response dated 17 November 2022 to s.174 notice of 27 October 2022, Part A question 31b, page 32.

customer demand was one of their key considerations when deciding which cloud provider to deploy their services on.[877]

6.74 In addition, the evidence received from cloud providers and ISVs for our interim report indicates that indirect network effects may favour more popular cloud providers, and so act as a barrier to entry and expansion. In particular:

a) Popular ISVs ([✂]) we have engaged with have only integrated their services with AWS, Microsoft or Google.[878] Some of these ISVs ([✂]) noted that they deployed their services on AWS first and expanded to Google and Microsoft one to four years later and mentioned that integrating with additional cloud providers requires material costs and time. For example, [✂] launched on AWS first, Azure two years later and on Google four years later. Similarly, [✂] launched on AWS first and on Azure and Google one year later.[879] Moreover, [✂] stated that achieving interoperability with an additional cloud infrastructure would involve significant costs.[880]
b) We understand that ISVs may also initiate open-source projects to develop tools that facilitate use of their proprietary services in the clouds where they are deployed (e.g. containers management). These tools will often be released on more popular clouds first (i.e. where the ISV service is available) and may not be fully functional on smaller cloud providers. For example, popular open-source services such as Terraform and Rancher offered support for AWS first and then expanded their products to Microsoft, Google and others. Moreover, we understand that Terraform may not offer the full set of functionalities on smaller clouds.
c) In line with the above, [✂] submitted that most challenges in cloud relate to the need to compete for developer attention to make the technology/solutions available on any cloud infrastructure. For example, for a period of time, VMware provided its market leading virtualisation technology solutions only on AWS.[881]
d) Our analysis of marketplaces broadly supports the above conclusion and suggests that the number of third-party services listed on the marketplaces of smaller cloud providers is likely to be significantly lower than the number of third-party services listed on the hyperscalers marketplaces.[882] For example, ([✂]) are only available on the marketplaces of the hyperscalers.[883]

---

[877] [✂] response dated [✂] to the s.174 dated [✂], Part A question [✂]. [✂] response dated [✂] to the s.174 dated [✂], Part A question [✂]. Ofcom / [✂] meeting, [✂].
[878] [✂] response dated [✂] to the s.174 dated [✂], Part A question [✂]. [✂] response dated [✂] to the s.174 dated [✂], Part A question [✂]. Ofcom / [✂] meeting, [✂]. [✂].
[879] [✂] response dated [✂] to the s.174 dated [✂], Part A question [✂]. [✂] response dated [✂] to the s.174 dated [✂], Part A question [✂].
[880] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[881] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[882] This is based on Ofcom analysis of responses to our statutory information requests to hyperscalers and of publicly available data on marketplaces. We recognise that third-party listed on marketplaces may include companies that are not ISVs (e.g. intermediaries, cloud consultants etc.). For this reason, we do not place much weight on the exact numbers of third-party services listed on specific marketplaces. However, we consider that the gap in third-party listings between hyperscalers and smaller providers' marketplaces may still provide a good qualitative indication that ISVs are generally more likely to list their services on hyperscalers' marketplaces.
[883] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂]. Ofcom / [✂] meeting, [✂].

6.75 In our interim report we noted that this type of indirect network effects may be mitigated by customers' ability to mix-and-match services on different clouds. If customers could mix-and-match services hosted on different clouds, they could easily access a range of services without having to purchase from a single cloud provider. This would make it easier for new or smaller cloud providers to attract customers and gain market traction. A hyperscaler [✂] said in response to our interim report that there is a high prevalence of multi-cloud usage, and as such, indirect network effects are weak.[884] Another hyperscaler respondent [✂] said that there is no evidence of network effects foreclosing competition, and even if such effects were presents, they would be mitigated by high prevalence of multi-cloud usage in the IT industry.[885] As set out in Section 4, whilst the use of multiple providers is relatively common, customers tend to use a secondary provider for use-cases that are more siloed and/or to back-up certain workloads. As we discuss below, customers may only procure a small number of services from a secondary provider and the majority of customer spend remains with their primary provider. The adoption of these multi-cloud architectures is therefore unlikely to mitigate network effects since ISVs would still have a greater incentive to deploy their services on clouds where the majority of workloads are being run (i.e. hyperscalers).[886] Therefore, we consider that indirect network effects are likely to persist unless customers are able to mix and match services on different clouds for more of their workloads as part of more integrated multi-cloud architectures.

6.76 In addition to this type of indirect network effects, cloud may exhibit two additional categories of network effects which may add to the barriers to entry and expansion for cloud providers. First, as set out earlier in this section, we consider that cloud exhibits a high degree of skills specialisation (i.e. the skills of cloud developers / engineers differ across clouds).[887] This may create some additional indirect network effects as the larger the pool of workforce proficient on a given cloud, the more likely a company will be to use that cloud (e.g. because it would be easier to find skilled operational engineers familiar with that provider). In turn, the larger the pool of customers using a cloud, the more people will choose to train to become proficient on that cloud.

6.77 Second, some customers have told us that in industries where cloud infrastructure services of separate users need to interact, customers may prefer to use cloud providers that are more popular amongst other users in their stakeholder group or supply chain. For example, [✂] told us that, amongst other considerations, their primary provider is AWS because most of their suppliers and customers were using it, which meant they would not be paying egress fees when exchanging data/content with them.[888] We also understand that choosing cloud providers that are popular within an industry may be a particular concern for start-ups who may wish to be acquired in the future. All things being equal, potential purchasers may look at start-ups who are using the same cloud provider more favourably, in order to minimise the necessary integration effort post-acquisition.

6.78 One hyperscaler, responding to our interim report [✂], said that there is no evidence that network effects in cloud give rise to competition concerns. It said that the entry or expansion of cloud providers is not hindered by scale advantages, whether due to network

[884] [✂] confidential response to the interim report, paragraph [✂].
[885] [✂].
[886] As demonstrated by shares of supply, where the portion of workloads that are run on smaller cloud providers is low compared to hyperscalers.
[887] See discussion in paragraphs 6.58-6.61 above for more detail.
[888] Ofcom / [✂] meeting, [✂].

effects or scale economies, and states that hundreds of new competitors have entered the UK cloud sector in recent years. It also said that in markets with significant network effects one would expect to see a correlation between the size of a competitor and their growth rate. They pointed to evidence that this has not been the case in the UK, stating that there has been high growth across all firm sizes.[889] However, we disagree with this assessment and consider that growth rates on their own are not informative about the presence of network effects. While the evidence presented by [✂] shows that a smaller provider [✂] experienced a similar growth rate to a hyperscaler [✂] between 2017 to 2022, [✂)]. As noted in Section 4, the joint share of hyperscalers continues to grow with the share of smaller providers getting smaller.

### Conclusion

6.79 In summary, we find that the need to build a broad product range raises barriers to entry and expansion for cloud providers. Our evidence suggests that the hyperscalers have developed their ecosystems to include the widest range of first- and third-party products with broad appeal. They also maintain a high rate of expansion of their product range through internal R&D and acquisitions. The hyperscalers benefit from significant economies of scale and scope, strong financial capabilities and access to technical expertise.

6.80 The evidence suggests there is scope to start with a narrower product range and expand more gradually. However, only Oracle and IBM have achieved a relatively wide product range, possibly due to benefits associated with having a wider software business (including access to technical skills and existing customer relationships). While they are closer to the hyperscalers, customer perception has not caught up, and they benefit less from network effects to attract ISVs and customers. Moreover, barriers to multi-cloud limit the ability for smaller cloud providers to gain material scale by competing for components of customer demand. As a result, smaller cloud providers have been unable to grow their market share materially, even in the case of Oracle whose product range is closest to that of the hyperscalers.

## Customer acquisition strategies

6.81 In this subsection we consider the importance of historical choice of cloud provider in determining customers' current choices and explore whether this creates barriers to entry and expansion for smaller providers. In terms of wider acquisition strategies, existing customer relationships in adjacent software markets and credit discounts appear to be important ways in which cloud providers attract new customers. We explore whether smaller cloud providers are able to leverage these strategies to the same extent as the largest providers.

### Importance of historical choices of cloud provider

6.82 Customers' historical choices and existing relationships with providers can significantly influence their future decisions on choice of provider. For example, one cloud provider ([✂]) suggested that while factors such as pricing and the range of services offered are

[889] [✂].

important in attracting customers to cloud providers, customers' historic relationships with cloud providers remains a "leading" factor in influencing their choice of cloud provider.[890]

6.83 AWS's and Microsoft's combined share of UK IaaS and PaaS revenues in 2022 stood at [✂]% [70-80%] and many customers will already have an established relationship with those providers. AWS is likely to have captured a material share of customers that were early adopters of cloud and Microsoft has built on its strong position in on-premises enterprise software (discussed further below). When new use-cases and workloads emerge, many customers are likely to take services from their chosen cloud provider because of the high switching costs and barriers to multi-cloud (discussed in Section 5). [✂].[891]

6.84 Our evidence indicates that in some cases there can be competition for new workloads that are sufficiently separate to a customer's existing workload hosted by their primary cloud provider (i.e. the more siloed end of the multi-cloud spectrum). Smaller cloud providers that entered the market later can compete for such new workloads and become the secondary provider for existing cloud customers (i.e. with AWS or Microsoft remaining as the primary provider).

6.85 However, at present it appears that the potential to build scale as a secondary provider may be limited. Given the many barriers to implementing an integrated multi-cloud architecture, secondary cloud providers are only able to compete for more siloed workloads. But siloed workloads may not be suitable for many of the use cases customers have.[892] Our evidence suggests that where customers have multi-cloud architectures, their spend is generally concentrated around a primary provider with only a small number of niche services being taken from a secondary provider.[893] By way of illustration, [✂] told us that [✂].[894] This is also supported by evidence collated from large customers, for example:

a) [✂].[895]
b) [✂].[896]
c) [✂].[897]
d) [✂].[898]
e) [✂].[899]
f) [✂].[900]

---

[890] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[891] [✂]; [✂].
[892] In a survey conducted by Public First, respondents that used more than one cloud were asked about the current level of integration between their different cloud providers and how important they considered the following types of integration: application integration, management integration, security integration and data integration. While the majority of respondents said that integration between different cloud platforms was 'very important' or 'somewhat important' for all types of integration, only 10% said their use of different cloud providers is currently largely integrated. The results of Public First's survey are available at: https://www.publicfirst.co.uk/files/CCIA_Survey.xlsx [accessed 19 September 2023]. The relevant questions are 49-50.
[893] For example, a [✂] report [✂] indicates that it is typical for multi-cloud organisations to concentrate 80% or more of their workloads with their primary strategic provider. [✂].
[894] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[895] [✂] response dated [✂] to our customer questionnaire.
[896] [✂].
[897] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].
[898] [✂] response dated [✂] to the s.174 notice dated [✂] question [✂].
[899] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].
[900] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].

6.86 We recognise there may be greater scope for smaller providers to compete for customers that are completely new to the cloud. According to data from IDC, spend on cloud services currently accounts for around 30% of total spend on IT services – both in the UK and globally.[901] IDC projects that the cloud deployment share of total IT spending will rise to 47% globally by 2027, with a slightly faster rise to 51% in the UK by 2027. Although the dataset does not include cloud usage on an individual customer-level basis, it is reasonable to assume that at least some of these workloads are from customers that do not currently purchase any cloud services – and therefore haven't yet chosen a primary cloud provider.[902]

6.87 However, these avenues for growth do not appear sufficient for smaller providers who have considerably lower shares than the market leaders, with Google the next largest provider at [✂]% [5-10%] in 2022 and Oracle remaining around [✂]% [0-5%] since they entered the market in 2016. The extent to which smaller providers can challenge for new customers is likely to be limited given the barriers set out earlier in this section. For example, customers value product range, so may be more inclined to go with a hyperscaler to access their broad product portfolios – or at least their reputation for having them. This may be particularly true for larger and more sophisticated customers, who tend to have more use cases. [✂].[903] This is likely to limit the potential for smaller providers to gain scale, given the vast sums large customers spend in the cloud. Our evidence indicates that revenues are concentrated among these largest customers, with workloads from approximately the top 1% of the hyperscalers' UK customers accounting for the majority of overall spend.[904]

## Importance of existing relationships in adjacent software markets

6.88 Beyond existing relationships for cloud services, our evidence indicates that existing relationships with providers for non-cloud services can also have an influence on customers' cloud choices. The market research found that 'existing relationships for other services' was one of the top ten reasons for choosing a cloud provider. Specifically, 'existing relationships for other services' was indicated in 23% of cases as an important factor for choosing a cloud provider (in 8% of cases it was the most important factor and in 15% of cases it was an important factor).[905] Of all factors that were chosen as 'most important' existing relationships received the third highest 'votes'.[906]

---

[901] Based on analysis of IDC Worldwide Black Book: Live Edition, July (V2 2023) Forecast (published July 2023). Total IT spending includes the following IDC technology categories from the Black Book publication: Infrastructure, Application Development & Deployment, Applications, System Infrastructure Software, Managed Services, Support Services, Project Oriented Services and Devices.

[902] As discussed in Section 3, while many organisations have already started migrating to the public cloud, there is still a proportion yet to begin this process.

[903] Data collected from responses to s.174 notices on the number of UK customers acquired in FY2020, FY2021, and FY2022 where each of these customers had a monthly spend greater than $100k on public cloud infrastructure services as at the last month of FY2022.

[904] [✂] Ofcom analysis of data from responses to our s.174 notices regarding UK customers of cloud infrastructure services. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]. This figure is disputed by one hyperscaler. [✂]. ([✂] response dated [✂] to our proposed use of information dated [✂]).

[905] Context Consulting research report, slide 69.

[906] After service quality (ranked first) and value for money (ranked second). Context Consulting research report, slide 69.

6.89 The ability for cloud providers to benefit from existing relationships in other markets may afford some cloud providers an important source of competitive advantage relative to others. In the interim report we said that Microsoft's position in traditional IT and SaaS makes its cloud services particularly attractive for midsize and large enterprises that are already using Microsoft's products on premises. We noted that Microsoft offers a range of enterprise software products and several of these have been estimated to have large market shares within their relevant product markets. For example, the CMA found that Microsoft has a share of 70-80% in the market for desktop operating systems.[907]

6.90 This is supported by the market research, which found that (amongst other reasons), 29% of respondents chose Microsoft Azure as their cloud provider, due to already having an existing relationship with Microsoft for other services and it was considered one of the top 6 reasons for choosing Microsoft.[908] This was higher than the average of 23% of participants across all providers saying that they chose their current provider due to already having an existing relationship for other services.[909] Some customers identify Azure as a natural choice when already using Microsoft for other services. One respondent [✂] agreed with our interim finding that Microsoft benefits from existing relationships with customers for its non-cloud services and noted that itself and other cloud providers do not have the legacy relationships to leverage in the same way.[910]

6.91 From our market research, we understand that some customers choose Azure due to the ease of integration it can offer with Microsoft's existing enterprise software products (e.g. Windows Server operating system, Microsoft 365 productivity software suite), with one customer in the market research citing that Azure's "integration with the other Microsoft systems is natural".[911] The ability for customers to integrate their cloud services with existing products can help simplify customers' IT management, and therefore, provide customers a degree of convenience.[912]

6.92 Microsoft's existing relationships with IT leaders and technical experts may also give several advantages to Microsoft. These existing relationships could help lower cloud migration costs for customers – customers that have an existing relationship with Microsoft for non-cloud products are already likely to have access to experts and staff with the relevant skillsets in using Microsoft. For example, ASOS explained that their choice of Microsoft for cloud was influenced by a combination of factors such as their internal skillset at the time, the range of PaaS services Microsoft offered and the engineering support it provided.[913] For these customers, choosing a different cloud provider would require them to retrain their existing staff or hire additional experts, which could bring about significant costs. Another advantage could be in relation to enterprise IT leaders endorsing Microsoft. Some research findings from [✂].[914]

6.93 Some respondents to the interim report noted that software licensing practices by certain legacy software vendors can disincentivise customers from using their existing software in

---

[907] CMA, 12 October 2022, Microsoft/Activision phase 1 decision, paragraph 260. [accessed 29 September 2023]
[908] Context Consulting research data tables, Q25.
[909] Ibid.
[910] [✂].
[911] Context Consulting research report, slide 57.
[912] Context Consulting research report, slide 56.
[913] Ofcom / ASOS meeting, 29 November 2022.
[914] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

another cloud environment.[915] Some customers also suggested there may be licensing advantages associated with using Microsoft Azure when choosing their cloud provider. For example, one customer ([✂]) explains that, having used Microsoft's Windows Server/SQL extensively on-premises, it decided to deploy the equivalent services mostly on Azure "for licensing reasons".[916] This customer ([✂]) suggested that it was advantageous to run Microsoft software on Microsoft's cloud, as they believe that it is cheaper to do so than to run the software on third-party clouds.[917] Another customer, [✂], told us that Microsoft provide "very strong commercial incentives" to use SQL Server on Azure and not on other cloud platforms; the customer gave examples of this, including examples of pricing incentives.[918] These pricing incentives could play a part in driving a customer's decision to choose Azure. Other licensing advantages could include familiarity with licensing model and established support making it a safer option. For example, [✂] recognise that existing relationships may provide a degree of familiarity to customers in the form of established support and known licensing models.[919] Some customers explained to us that licensing advantages associated with using Microsoft Azure are one factor they consider when choosing their cloud provider out of a variety of factors driving the customers' decisions.[920]

6.94 Mid-scale cloud providers, such as IBM and Oracle are also known to provide on-premises services (e.g. Oracle databases), whilst VMware is known for its private cloud platform. Therefore, these suppliers may also benefit from their position in adjacent software markets and private cloud platforms. For example:

a) The market research shows that for 23% of Oracle cloud customers surveyed (note low base of 100 Oracle customers) and 26% of IBM cloud customers surveyed indicated that 'existing relationships for other services' was one of the factors they considered when choosing their cloud provider.[921]
b) In our engagement with large customers, [✂] explained that it is currently considering the use of Oracle for its (on premises) Oracle workloads due to its existing relationship with Oracle for non-cloud services.[922]
c) Oracle has indicated that it generally considers that it is easier to move a customer from on-premises to the cloud, where the customer has an existing, on-premises relationship with that cloud provider.[923]
d) Similarly, IBM recognise that "many clients will choose a cloud provider with whom they have existing relationships, which has helped the growth of IBM's Cloud Platform".[924]

---

[915] Computer & Communications Industry Association response to the Interim Report, question 5.1. Computer & Communications Industry Association (ofcom.org.uk) [accessed 21 September 2023]; and [✂] confidential response to the Interim Report, question [✂].
[916] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[917] [✂] response dated [✂] to our follow-up questions by email dated [✂] concerning the s.174 notice dated [✂], question [✂].
[918] [✂] response dated [✂] to our follow-up questions by email dated [✂] concerning the Ofcom / [✂] meeting, [✂].
[919] [✂] response dated [✂] to the s.174 notice dated [✂], page [✂], question [✂].
[920] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] via email on [✂]; and Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] via email on [✂].
[921] Context Consulting research data tables, Q25.
[922] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[923] Oracle response dated 20 March 2023 to our proposed use of information dated 14 March 2023.
[924] IBM response dated 6 December 2022 to the s.174 notice dated 25 October 2022, Part A, question 7.

6.95 For these cloud providers, their position in an adjacent software market therefore is likely to have contributed to their ability to enter the market and gain some market share. However, for the reasons already set out above (less extensive product offerings, etc.), mid-scale cloud providers are likely to find it more difficult to win new customers, especially those with existing relationships with Microsoft.

6.96 There are also some suggestions that AWS and Google may benefit from existing relationships. One cloud provider ([✂]) said that many of AWS's customers today are those that initially started up on AWS, using Amazon's platform for the sale and distribution of their products.[925] In one case, AWS appears to have been able to extend its contract by guaranteeing a customer access to one of Amazon's streaming devices. For example, it is reported that WarnerMedia was able to launch its streaming service (HBO Max) onto Amazon's Fire TV device, only after agreeing to extend its contract with AWS.[926] Another cloud provider ([✂]) argues that Google leverages its existing relationships with Chief Marketing Officers,[927] and continue to "bundle cloud offers with Advertising and Google Workspace" to attract new customers.[928] [✂].[929] However, our evidence on this is only anecdotal, and the market research suggests that it is not considered one of the top 6 reasons for choosing AWS and Google.[930] Furthermore, the market research finds Google and AWS benefit from relationships in other services to a lesser degree than Microsoft.[931]

## Cloud credits

6.97 Cloud credits appear to be an important feature of cloud providers' acquisition strategies for customers that do not already have an existing relationship with any cloud provider, such as start-ups. The cost of credits can be substantial and could pose a barrier to entry and expansion for the smallest cloud providers.

6.98 In Section 4, we explored how cloud credits work and the cloud providers' rationale for offering these. The credits offered by cloud providers can typically be spent on services across IaaS and PaaS. We noted that most providers offer free credits for customers opening an account with them for the first time. These credits tend to be of low monetary value and are generally comparable across both the hyperscalers and smaller cloud providers. For example, both Microsoft and IBM offer $200 for 1 month.[932]

6.99 Credits offered to start-ups and scale-ups are of much higher monetary value than the credit programs generally offered via cloud providers' websites to new customers. Cloud providers also have partnerships with venture capital firms, where they can reach start-ups/scale-ups and offer exclusive benefits to the venture capital firm's portfolio of companies. There is variability across cloud providers, both in terms of the monetary value and time limits associated with their credit offerings. For example:

---

[925] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[926] The Information, 2021, WarnerMedia Extended AWS Deal to Win Key HBO Max Concession [accessed 19 September 2023].
[927] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[928] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[929] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[930] Context Consulting research report slide 70; and Context Consulting research data tables Q25.
[931] The market research found that 29% of Microsoft customers listed 'having an existing relationship for other services' as a reason for choosing Microsoft as its cloud provider, compared to 14% of AWS customers and 22% Google customers. Context Consulting research data tables, Q25.
[932] Azure website, Azure free account; and IBM website, IBM Cloud free tier [accessed 19 September 2023].

a) The hyperscalers offer substantial monetary credits to start ups and scale ups. For example, AWS offers up to $100k for the first year, Google offers up to $100k for each year for the first two years, so a total of $200k, and Microsoft offers up to $150k for the first year.[933]
b) IBM's credits to start-ups and scale-ups programme offers eligible start-ups up to $1000 per month for 12 months or $3000 per month for 6 months.[934] IBM offers credits of a monetary value that is more comparable to the hyperscalers' start-up/scale-up programmes via partnerships with venture capital firms. For example, it offers up to $120,000 in IBM Cloud credits to eligible portfolio companies at DSW Ventures and Aurelia Ventures.[935]
c) Oracle's website does not currently specify the monetary value of cloud credits or discounts it offers to start-ups. Oracle told us that it offers $300 in cloud credits to anyone, including start-ups. In addition, Oracle offers, on request, $500 in cloud credits to developers.[936] There is also some evidence to suggest that Oracle offers a 70% discount on Oracle cloud for start-ups, but this offer may have been withdrawn, as it's not currently evident on its live website.[937]
d) OVHcloud's start-up programme offers up to 100k euros (or equivalent local currency) for the first year.[938]
e) Scaleway's start-up programme offers up to 36k euros for the first year.[939]

6.100 We also have some evidence of cloud credits being used to lower the cost of migration from on-premises: two customers ([✂] and [✂]) told us that they each received substantial cloud credits of $[✂]m and $[✂]m respectively, over a 5-year term from Google to assist with services such as data migrations.[940] [✂] told us that Oracle offered cloud credits as an incentive to migrate workloads from Oracle on-premises solutions to Oracle's cloud.[941] These types of credits lower the costs of migrating to the cloud from on-premises IT and therefore could be an important strategy for gaining new cloud customers.

6.101 One cloud provider in their response to our consultation ([✂]) said that evidence of smaller cloud providers offering credits as part of start-up programmes suggests that there are no particular barriers to offering credits.[942] Indeed, our evidence shows that some smaller providers are able to offer credits of a similar magnitude to the hyperscalers – including OVHcloud as part of its start-up programme, or IBM via partnerships with venture capital firms. However, this is not universally the case and some small-scale cloud providers ([✂])

---

[933] AWS website, AWS activate; Azure website, Unlocking Azure credits as your start-up grows; and Google website, Google for start-ups cloud program [accessed 19 September 2023].
[934] IBM website, The start-up with IBM program [accessed 19 September 2023].
[935] Aurelia Ventures website, IBM Cloud credits; and DSW Ventures website, IBM [accessed 19 September 2023].
[936] Oracle email to Ofcom, dated 13 September 2023.
[937] The Dutch competition authority (ACM) reported these figures from Oracle's website which it last accessed in April 2022. ACM, September 2022, Market Study Cloud services, page 45. Furthermore, this is supported by a pdf document on Oracle's website, although this may be an archived document. Oracle website, Oracle for Startups [accessed 19 September 2023].
[938] OVHcloud website, Startup programme FAQs [accessed 19 September 2023].
[939] Scaleway website, Why choose Scaleway? [accessed 19 September 2023].
[940] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂] question [✂].
[941] [✂] response dated [✂] to the s.174 notice dated [✂] question [✂].
[942] [✂].

told us that they are unable to match the credits offered by the hyperscalers.[943] This is supported by The Sustainable Digital Infrastructure Alliance (SDIA), who said that most cloud providers – especially new entrants – cannot afford to offer the same types of credits as the hyperscalers.[944]

6.102 This is consistent with evidence that the hyperscalers issue the highest amounts of credits, both in absolute terms and relative to their cloud revenue. The hyperscalers issued substantial credits to UK customers in 2022: AWS issued $[✂] worth of credits (approximately [✂]% of its public cloud revenue) and Google issued $[✂] worth of credits (approximately [✂]% of its public cloud revenue).[945] Whilst Microsoft was unable to provide data on cloud credits issued to UK/Global customers, it did provide an estimate of cloud credits redeemed by UK customers of $[✂].[946] Mid-scale providers issued lower amounts of credits than the hyperscalers, though still relatively substantial: IBM issued $[✂] worth of credits (approximately [✂]% of its public cloud revenue) and Oracle issued $[✂] worth of credits (approximately [✂]% of its public cloud revenue).[947] Furthermore, AWS's credit issuance exceeds the UK public cloud revenue of some smaller providers (e.g. [✂] UK public cloud revenue of $[✂]).[948]

6.103 In comparison, small-scale providers appear less able to offer credits on the same scale as the hyperscalers. [✂] explained that its credits are much more limited in comparison to the hyperscalers' in monetary terms [✂] and the total number of companies it can offer credits to is [✂].[949]

6.104 This may be because small-scale providers are less able to monetise customers acquired through credit programmes than the hyperscalers. The French competition authority, *Autorité de la concurrence,* cited evidence from one hyperscaler in its market study that the future revenues generated by cloud credits are far greater than their initial cost, and that they are able to generate a positive return on investment within three years or shorter.[950] In comparison, small cloud providers may be less able to make such a commercial policy profitable in a similarly short space of time due to having more limited product portfolios. The hyperscalers offer a broad range of services, so once a start-up/scale-up is acquired there are many opportunities for the hyperscalers to upsell and expand the number of services purchased by the customer. In comparison, small cloud providers offer a narrower range of services, so the time required for each customer to become profitable may be longer. This may make it less feasible financially for small cloud providers to acquire

[943] Ofcom / [✂] meeting, [✂] and Ofcom / [✂] meeting, [✂].
[944] The SDIA are a non-profit think tank focused on digital infrastructure. SDIA response to the interim report: Vision for a Sustainable, Federated European Cloud, page 4.
[945] Note that the total amount of credits issued may not have been redeemed by customers during that period. AWS response dated 31 March 2023 to the s.174 notice dated 24 October 2022, Part B question 8; Google response dated 31 March 2023 to the s.174 notice dated 26 October 2022, Part B question 8.
[946] Microsoft response dated 20 July 2023 to our follow-up email dated 10 May 2023 concerning the s.174 notice dated 21 October 2022, Part B question 8.
[947] IBM response dated 4 April 2023 to our s.174 notice dated 25 October 2022, Part B question 8; Oracle response dated 30 March 2023 to the s.174 notice dated 31 October 2022, Part B question 8; Oracle response dated 2 June 2023 to our follow-up email dated 10 May 2023 concerning the s.174 notice dated 31 October 2022, Part B question 8.
[948] AWS response dated 31 March 2023 to the s.174 notice dated 24 October 2022, Part B question 8; [✂] response dated [✂] to the s.174 notice [✂], Part B question [✂].
[949] Ofcom / [✂] meeting, [✂].
[950] *Autorité de la concurrence*, 2023. Opinion 23-A-08 of June 29, 2023 on competition in the cloud sector [accessed 20 September 2023].

customers through costly credit programmes.[951] Indeed, one small-scale cloud provider ([✂]) explained that it cannot afford/absorb the losses incurred during these credit periods.[952] There is also a risk that a start-up/scale-up may go bankrupt and therefore the investment made in terms of credits will not be recouped. This risk can be reduced for the hyperscalers who may be able to diversify their portfolio of risk more effectively. The hyperscalers may be better able to attract start-ups/scale-ups across a variety of different industries because they offer a broad range of products/services with wide appeal across industries.

6.105 Furthermore, customers' familiarity with the hyperscalers and a perception that they represent the safe option in cloud may also be relevant factors.[953] Our market research found that few customers had considered using providers outside of the hyperscalers, with some customers lacking any awareness of small-scale cloud providers.[954] There may be various other factors specifically influencing start-ups' decisions on choice of cloud provider, e.g. in industries where acquisitions are common, there may be a preference to use the same cloud provider as future potential acquiring companies.

### Conclusion

6.106 Overall, we have found that the hyperscalers (Microsoft and AWS in particular) are better able to expand their sales relative to smaller cloud providers.

6.107 Many customers already have an established relationship with AWS or Microsoft in cloud and barriers to switching and multi-cloud imply they are likely to continue to use them in future. For some of these customers, Google, IBM and Oracle can compete to become the secondary provider for workloads that are sufficiently separate to a customer's existing cloud usage. But the potential to gain scale in this way may be limited.

6.108 Our evidence also suggests that Microsoft may benefit from its existing relationships with customers for its non-cloud services, e.g. enterprise software and on-premises services. We acknowledge that smaller providers, such as Oracle and IBM, can also benefit from their position in adjacent software markets and this is likely to have contributed to their ability to enter the market and gain a small market share.

6.109 Cloud credits are an important acquisition strategy for attracting new customers where the hyperscalers do not already have an existing relationship. We find that the hyperscalers offer the highest amounts of credits, both in absolute terms and as a proportion of revenue, followed by mid-scale providers. While small-scale providers can also offer credits, their lack of financial resources mean that they are unlikely able to match the hyperscalers – for example, [✂] told us its credits are much more limited in terms of monetary value and the total number of companies it can offer credits to.

## Conclusion on barriers to entry and expansion

6.110 The factors we identify in this section combine to pose material barriers to entry and expansion in the provision of cloud infrastructure services. In particular in relation to the

[951] [✂]. [✂] confidential response to the interim report, [✂].
[952] Ofcom / [✂] meeting, [✂].
[953] Context Consulting research report, slides 64-65.
[954] Context Consulting research report, slides 64-65.

ability of cloud providers to develop a broad ecosystems of services powered by a global network of data centres.

6.111 Our evidence suggests that AWS and Microsoft have faced the fewest barriers, which has allowed them to materially pull ahead of other cloud providers in terms of their customer base and associated scale. Our evidence points to some differentiation in the factors that have encouraged customers to take-up their respective ecosystems:

a) AWS is likely to benefit from a first-mover advantage, which has allowed it to gradually build its customer base and phase its investment in product range and infrastructure ahead of others entering the market. As a result, it is recognised as offering the broadest range of first- and third-party products, and likely benefits from economies of scale from its global network of large-scale data centres. It would appear to benefit most from network effects to attract ISVs and new customers into its ecosystem.
b) Microsoft has caught-up with AWS in terms of its broad product range and network of large-scale data centres. It rivals AWS's large established customer base, which means it is likely to benefit from network effects as it attracts ISVs and new customers into its ecosystem. In contrast to AWS, Microsoft is likely to benefit from its leadership position in adjacent software markets. Our evidence suggests that customers using Microsoft's business enterprise suite can more easily integrate these services with Azure and are likely to have access to staff with the relevant skillsets that can be more easily transferred to Azure.

6.112 We recognise that some cloud providers have entered in recent years. However, their scale remains materially below that of the market leaders and, based on IDC data, there has been no significant cloud provider entry in the UK since 2018 (see Section 3). This is likely to reflect the fact that smaller providers have faced more barriers to grow scale in cloud infrastructure services, reducing the effective constraint that they exert on the market leaders:

a) Google has been gaining customers, having built-out a broad portfolio of cloud services and with particular strengths in data analytics. Google's position in cloud is likely aided by its large non-cloud businesses acting as an anchor tenant and capabilities developed in adjacent digital markets. However, its share remains far behind the market leaders. Due to its smaller scale, Google is less able to benefit from network effects. For large customers already established with AWS and Microsoft, barriers to switching and multi-clouding imply that Google may be restricted to compete to become a secondary provider for some siloed workloads only. Google is unlikely to benefit from existing relationships with enterprise customers to the same extent as Microsoft.
b) IBM and Oracle are likely to benefit from their position in adjacent software markets to migrate existing customers into their public clouds, and to draw on existing software engineering skills to build-out their range of PaaS products. However, barriers to multi-cloud and switching are likely to limit their ability to gain scale by competing for components of customer demand that can be served by their narrower product range. Their small customer base and lack of anchor tenant implies they are likely to have cost disadvantages. Moreover, their lack of scale is also likely to reduce network effects. For example, compared to the hyperscalers, there will be a smaller number of ISV services available on their clouds and a smaller pool of engineers trained to work with these providers' clouds. This suggests they pose only a limited constraint on AWS and Microsoft when customers choose their cloud ecosystem provider.

c) Small-scale providers are likely to be affected by all the barriers to entry and expansion we have identified. They likely pose a negligible constraint on the market leaders as they appear to challenge for a narrower set of potential cloud infrastructure customers and tend to have more specialised offerings.

6.113 In contrast, the barriers to entry and expansion are lower for the supply of individual products in PaaS, as ISVs can build on the physical infrastructure of other cloud providers. Some ISVs appear able to compete head-on with the hyperscalers for specific workloads, although none of these can challenge the hyperscalers in terms of range. As a result, even though there are many ISVs, they cannot compete with cloud providers on an equal footing. Additionally, the potential reliance of ISVs on the hyperscalers may raise other potential issues, which we discuss in Section 7.

# 7. Hyperscalers' relationship with ISVs

7.1 In this section, we discuss the different types of relationships that hyperscalers may have with ISVs and how these relationships may impact competition in the cloud market.

7.2 Since publishing our interim report, we have only received additional evidence from a small number of ISVs. This section updates our analysis and concerns relating to hyperscalers' relationships with ISVs.

## Hyperscalers' relationships with ISVs could impact competition

7.3 Hyperscalers interact with ISVs in a number of different ways:

a) Hyperscalers can act as input suppliers to ISVs. Specifically, hyperscalers provide cloud infrastructure services and, in turn, ISVs rely on these services to develop and run their own cloud services which may complement or compete with those offered by the hyperscalers.
b) Hyperscalers can also act as distributors of ISVs' services and provide ISVs with a route to market. This could be through directly selling ISVs' services, offering ISVs a platform through which to sell their services (such as a marketplace) or access to customers.
c) Since hyperscalers operate and sell services across all layers of the cloud stack, they can also compete directly with ISVs, offering services at the same layer of the cloud stack as ISVs.

7.4 As a result, the hyperscalers may in some cases have a dual role. On the one hand, they provide cloud infrastructure services to ISVs and act as a distributor for their services. On the other hand, they offer cloud services that compete with those of ISVs which are hosted on their clouds. This may create a potential conflict of interest and give the hyperscalers the opportunity to provide their own services with an advantage over ISVs' competing services, ultimately increasing barriers to entry and expansion for ISVs.

7.5 In principle, the hyperscalers might limit entry and expansion of competing ISVs in two ways:

a) In their role as suppliers of cloud infrastructure services, hyperscalers may deny, restrict or increase costs of access to the cloud infrastructure services that ISVs need to effectively run their services. This could be done through technical mechanisms (such as not providing the necessary public APIs) or commercial mechanisms (such as raising prices for ISVs' use of the hyperscalers' cloud infrastructure services). As a result, ISVs might see a reduction in the quality or ease of use of their services, compared to equivalent services which are offered by the hyperscalers and are not subject to these limitations.
b) In their role as distributors (particularly via marketplaces), hyperscalers might be able to raise ISVs' costs (for example, by increasing marketplace commission fees), self-preference their own services over ISVs' (e.g. by making their own services more prominent), or gain information advantages (such as an understanding of which products perform well).

7.6 In the remainder of this section, we separately consider hyperscalers' roles as input suppliers to ISVs and as distributors of ISVs' services. For each of those roles, we assess the extent to which: i) ISVs are reliant on hyperscalers; and ii) hyperscalers are engaging in any practices that may increase barriers to entry and expansion for ISVs.

# Hyperscalers as suppliers of cloud infrastructure services to ISVs

## ISVs deploy their services on the hyperscalers' clouds to access their customer base

7.7 Evidence from ISVs indicates that they typically deploy their services across the three hyperscalers, at least initially, to access their respective user bases. For example, three ISVs ([✂]), submitted that expected customer demand was one of the key considerations when deciding which cloud to deploy their services on.[955] [✂] also noted that they deployed their services on AWS first and expanded to Microsoft and Google one to four years later.

7.8 In principle, it would be possible for ISVs to start self-supplying their own cloud infrastructure and operate as a cloud provider. However, this would be challenging, given the material barriers to entry and expansion set out in Section 6. [✂] submitted that, whilst theoretically possible, self-supplying all or some of the necessary cloud infrastructure services or hardware would be a significant investment and undertaking.[956] Similarly, [✂] submitted that it would be difficult to self-supply cloud infrastructure in all of the hyperscalers geographical regions and deliver its service from outside the hyperscaler regions, due to potential latency issues which are critical for its [✂] applications.[957]

7.9 In addition, an ISV hosting its services in one cloud might offer its services to customers of a different cloud (i.e. a customer using AWS might be able to use an ISV service hosted on Google). However, this would involve integrating services from different clouds which can be challenging given the high barriers to adopt an integrated multi-cloud architecture (as discussed in Section 5).

7.10 In summary, ISVs have little alternative but to deploy their services on the hyperscalers' clouds to gain access to the hyperscalers' customers. We go on to discuss how the reliance by ISVs on hyperscalers influences how competition plays out, beginning with how ISVs technically interoperate with hyperscalers' clouds.

## ISVs can access the hyperscalers' cloud infrastructure services, but we have heard some concerns that hyperscalers may favour their own services

7.11 The hyperscalers submitted that they provide the technical information and support needed by any third parties (including ISVs) to fully interoperate with their cloud infrastructure services or hardware. This includes, in particular, documentation about the services available

[955] Ofcom / [✂] meeting, [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[956] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[957] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].

and information on how third parties can interoperate with them (e.g. SDKs, APIs and protocols). More specifically:

a) AWS said that it makes many of its SDKs and APIs publicly available under open-source licences, so that customers and third parties may freely use, modify and distribute them without restrictions. In addition, AWS mentioned that its services support various standard protocols to make it easier for third parties to enable communications and interactions between services in a common way.[958]
b) Microsoft said it makes extensive information available to developers about the services available in Azure and how any third parties can access that functionality. We understand that, similar to the other hyperscalers, Microsoft APIs are easily accessible via command line tools and SDKs, which exist for many popular programming languages, as well as other common technical routes.[959]
c) Google said that in order to facilitate interoperability, it offers Google Cloud APIs (i.e. interfaces to Google's cloud services) which are publicly available, including to third-party suppliers of cloud services. Third parties can access Google Cloud APIs in many popular programming languages via a variety of technical routes.[960]

7.12 The hyperscalers also submitted that they do not restrict access to their clouds and generally provide third parties with the same set of features and functionalities they make available to their own services. Specifically:

a) AWS said it does not limit interoperability either contractually or technically but rather strive to enable interoperability with the systems and services of other IT providers to serve AWS's customers.[961] AWS noted that the level of interoperability of AWS services with third-party offerings depends on a variety of factors, including the third-party offering's support for standard protocols (providers whose offerings support more protocols may be more interoperable with AWS); network connectivity between the AWS data centre and the third party's servers (longer distances and lower quality networks may suffer increased latency); and the adoption of common open-source software components by AWS and third parties.[962] AWS also said that if customers using AWS's services could not easily incorporate third-party software into their solutions, AWS would be unable to attract new customers or retain customers seeking to utilise third-party solutions.[963]
b) Microsoft said that it makes the same features of Azure available to its customers and other third-party suppliers without discriminating based on whether the entity is a provider of cloud services. This includes designing Azure to allow third-party PaaS services to have the same opportunity as Microsoft's own services. Microsoft said that making Azure an attractive platform on which third-party services providers can deploy

[958] AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, Part A question 31, paragraphs 31.6 and 31.7; and [✂].
[959] Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, Part A question 31; and Microsoft response to the interim report dated 31 May 2023, paragraphs 182(i), 208.
[960] Google response dated 23 November 2022 to the s.174 notice dated 26 October 2022, Part A question 32a, page 59-60; and [✂].
[961] AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, Part A question 31b, paragraph 31.11.
[962] AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, Part A question 31a(iii).
[963] [✂].

their solutions is critical to its ability to compete.[964] Finally, Microsoft said that scale benefits are passed onto PaaS providers through equal access to underlying infrastructure, and as a result PaaS-only providers compete on an equal footing with the vertically integrated providers.[965]

c) Google said that interoperability with third-party services is at the heart of Google's cloud services proposition. This means that, wherever it is technically possible to do so, Google Cloud enables full interoperability between its own cloud platform and the services of third-party suppliers of cloud services.[966]

7.13 The feedback received from ISVs confirms they are typically able to access the minimum set of services and functionalities to run their services on the hyperscalers' clouds. For example:

a) [✂] submitted that to date, it has not encountered any circumstances that have prevented its service from achieving an ideal level of interoperability with the public cloud infrastructures on which its service operates.[967]

b) [✂] submitted that all the hyperscalers open their APIs sufficiently to allow it to integrate its software ([✂]) and achieve similar performance to the hyperscalers first-party services in relation to core software functionality.[968]

c) [✂] submitted that it has been able to develop a very tight integration with the hyperscalers clouds to enable automatic deployment with the click of a button.[969]

7.14 However, some ISVs gave examples of cases where hyperscalers only allow access to certain cloud functionalities for the benefit of their first-party products, thereby negatively affecting ease of use by customers of ISVs' services. In particular:

a) [✂] said that customers are typically required to complete additional steps to set up and manage ISVs' services compared to first-party services. It highlighted specific limitations and lack of access to certain functionalities of the hosting hyperscaler cloud, limiting its ability to integrate its service. This increases friction and discourages take-up of ISVs.[970]

b) [✂] said that ISVs' services would appear less embedded into AWS's and Microsoft's clouds which may affect the overall user experience. For example, unlike AWS's first-party [✂] services ([e.g. ✂]) where the underlying infrastructure with all its complexities is managed by AWS and hidden away from the user, a customer has much more visibility into how [✂] services are deployed which makes them appear less polished. The current level of interoperability with AWS does not allow [✂] to achieve the same level of integration which puts them at a disadvantage compared to native first-party services.[971]

c) [✂] said that keeping its cloud services up-to-date, such that they can take advantage of new functionalities of the underlying first-party cloud services introduced by the cloud provider, can require very substantial technical work running into months and even a year. The ISV suggested that this issue may be mitigated by requiring cloud providers to

---

[964] Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, Part A question 31a.
[965] Microsoft response to the interim report dated 31 May 2023, paragraph 75.
[966] Google response dated 23 November 2022 to the s.174 notice dated 26 October 2022, Part A question 32.
[967] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[968] Ofcom / [✂] meeting, [✂].
[969] Ofcom / [✂] meeting, [✂].
[970] Ofcom / [✂] meeting, [✂]; Ofcom / [✂] meeting; [✂] and Ofcom / [✂] meeting, [✂].
[971] Ofcom / [✂] meeting, [✂]; Ofcom / [✂] meeting, [✂].

use standardised APIs and to develop any forked versions of open-source software in a more transparent way.[972] Microsoft acknowledged this issue in its response to the interim report. When comparing third-party cloud services to its own cloud services, Microsoft said that an ISV or third-party cloud provider may not be able to fully test and update their applications until sometime after Microsoft's new API release. As a result, there may be a reduction in interoperability between applications updated at different times, leading to lower performance or product quality (i.e. certain features do not work correctly).[973] In addition, [✂].[974] We note that [✂].[975]

7.15 Some other respondents also commented on ISVs' interoperability with the hosting clouds. BT Group said that cloud services offered by cloud providers should be 'unbundled' or 'separated' into their individual elements to allow third-party cloud services to interoperate with those individual elements.[976] The Federation of Communications Services said that, in principle, it would not support hyperscalers being able to limit entry and expansion of competing ISVs (for example, by restricting APIs or inappropriately raising ISVs' costs) and that it wants to see a very competitive ISV market, encouraging new and wide-ranging functionality and interoperability improvements, therefore suggesting that current levels of interoperability are insufficient.[977]

7.16 In addition, two ISVs ([✂]) said that the hyperscalers may also be using other technical mechanisms to exploit their position as cloud providers and unfairly compete with ISVs. These ISVs explained that, in some cases, the hyperscalers build first-party services on open standards and open APIs, but introduce proprietary features that require customers wanting to (integrate and) switch to rewrite a significant portion of their code. Customers may be induced to take-up these first-party cloud services (instead of the original open-source ones) because, as noted above, AWS and Microsoft make them more integrated and easy to use, or because they may not fully disclose the added proprietary features such that some customers may not be aware these are different from the original open-source version. More specifically:

a) An ISV ([✂]) explained that AWS and Microsoft have developed proprietary [✂] services which they sell as compatible with the [✂] software. Customers may take up these services because they are more visible, easy to use or because customers are under the impression – driven by a lack of transparency from these hyperscalers – that these adaptations are fully compatible with (or even a version of) the [✂] service. However, according to [✂], the compatibility of these first-party cloud services is limited, meaning that a customer using such [✂] compatible services would need to considerably reconfigure their application if they wanted to switch to [✂].[978] [✂].[979]
b) Similarly, another ISV ([✂]) explained that the hyperscalers and other large platform vendors unfairly compete with best-of-breed innovators within their ecosystems, by

[972] [✂] confidential response to the interim report, page [✂]. We have also heard this from suppliers of professional services. See Sections 5 and 6.
[973] We understand that Microsoft attributes these limitations to the natural result of innovation from cloud providers. Microsoft response to the interim report dated 31 May 2023, paragraph 192.
[974] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[975] Ofcom / [✂] meeting, [✂].
[976] BT Group response to the interim report, page 4.
[977] Federation of Communication Services response to the interim report, question 5.3.
[978] Ofcom / [✂] meeting, [✂].
[979] [✂].

building their first-party services on open standards and open APIs but introducing proprietary features.[980] This makes it more difficult for customers to switch from first-party to third-party software within the hyperscalers' ecosystems as customers would need to rewrite a significant portion of their code.[981] We note that this includes customers wishing to switch to cloud-agnostic ISVs' services to facilitate integration of multiple clouds and switching between clouds.

7.17 We have also received evidence that the hyperscalers may have an advantage as cloud providers can offer bundles of discounted services that ISVs are not able to match.[982] For example, an ISV ([✂]) told us that the main expansion challenge it has faced is the aggressive pricing the hyperscalers practice on comparable first-party services,[983] since they are able to offset aggressive pricing on one product with the additional revenue generated by other products in their ecosystems.[984] Another ISV ([✂]) said that adoption of its service has been limited, because procuring multiple services from a cloud provider can be cheaper and easier. In particular, when using its services in combination with those of a cloud provider, customers have to pay the fee charged by the ISV, in addition to the one charged by the cloud provider. Conversely, when customers source all their services from a single cloud provider, they would only pay a single fee to that cloud provider.[985]

7.18 Google disagreed. It said that, while it offers a variety of discounts, it does not offer any predetermined bundles or packages of its cloud infrastructure services, or any fixed bundles or packages that combine its cloud infrastructure services with other non-cloud Google products/services.[986]

7.19 SDIA generally agreed with the ISVs concerns above. It said that the marketplace invites further packaging/bundling of services (selling software products which include digital resources/infrastructure/IaaS of the cloud provider) which creates further barriers that make it difficult for national and regional cloud providers to compete with the hyperscalers.[987]

7.20 Lastly, two ISVs ([✂]) indicated that they are generally wary of potential future risks of the hyperscalers self-preferencing their own services, including through technical restrictions. However, they did not present any evidence of these risks currently materialising.[988] In their response to our interim report, an individual, Priyank Chandra, said that, once ISVs create value in the market, hyperscalers become primary competitors to ISVs, since they can provide more integrated services.[989]

---

[980] The introduction of proprietary features can be pro-competitive where such features are the result of innovation, e.g. to introduce new or improved functionality. However, it can also be anti-competitive where such features are not the result of innovation (i.e. there is no new or improved functionality) but aim to technically differentiate the product such that ISVs' cloud services appear less integrated.
[981] Ofcom / [✂] meeting, [✂].
[982] Relatedly, another obstacle to the take-up by customers of the services offered by ISVs is that spend on ISVs' services may only partially contribute to customers' committed spend with hyperscalers. See Section 5 for additional details.
[983] [✂].
[984] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[985] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[986] [✂].
[987] SDIA response to the interim report, page 3, question 5.3.
[988] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂]. [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice [✂], Part A question [✂]. [✂].
[989] Priyank Chandra response to the interim report, question 5.3.

## Conclusion on hyperscalers as suppliers of cloud infrastructure services to ISVs

7.21 Overall, the evidence we have received during this market study suggests that ISVs materially rely on the hyperscalers' cloud infrastructure to access customers on the hyperscalers' ecosystems. The feedback received from ISVs suggests that hyperscalers are currently providing the technical information and support needed by ISVs to host and run their services on the hyperscalers' clouds. However, we have also identified how the hyperscalers might be able to discourage the take-up of competing services developed by ISVs e.g. through technical restrictions.[990] ISVs have presented some evidence of specific cases where hyperscalers may be doing this, though this is disputed by hyperscalers and for the purpose of this study we have not found it necessary to reach a definitive view on this.

7.22 Our view is that the hyperscalers still appear to have strong incentives to attract ISVs onto their clouds and expand the range of services they can offer to customers and drive the usage of their underlying infrastructure. This means they currently may have limited incentives to systematically restrict the availability or quality of the ISVs' services they host on their clouds.

7.23 However, our conclusions are based on responses received from, and meetings with, a relatively small number of ISVs during the market study. We also acknowledge that the risk of the hyperscalers restricting the availability or quality of ISVs' services on their clouds may increase if the incentives for the hyperscalers to draw ISVs onto their ecosystem were to decrease in future. We discuss this further in Section 8.

# Hyperscalers as distributors of ISVs' services

7.24 Hyperscalers distribute ISVs' services either directly (by selling the services of ISVs directly to business customers) or indirectly (through the marketplaces they operate or by providing support to ISVs looking to sell services which run on their clouds).

7.25 Hyperscalers told us that it is rare for them to directly sell the services of ISVs to business customers.[991] However, as explained in Section 4, the hyperscalers offer co-sell support to ISVs and each hyperscaler offers a marketplace (an online platform), where providers can offer services to customers which run on the underlying infrastructure of the marketplace provider.[992]

7.26 Marketplaces can allow ISVs to access a broader range of customers and so can support their ability to compete with the hyperscalers. However, there are risks for ISVs if they become too reliant on marketplaces to access customers. If the hyperscalers could control the entry point to reach customers, this could provide them with the ability to exploit ISVs and ultimately distort competition. Such practices would act as a barrier to entry and

[990] These may include hyperscalers technically restricting ISVs' access to existing or new proprietary features or services.
[991] [✂]. See [✂] response dated [✂] to s.174 notice dated [✂], Part A question [✂]; [✂] response dated [✂] to s.174 notice dated [✂], Part A question [✂]; [✂] response dated [✂] to s.174 notice dated [✂], Part A question [✂].
[992] We understand that a number of other smaller suppliers of cloud services operate marketplaces. These include IBM (Cloud Catalog), Oracle, OVHcloud and Salesforce (AppExchange). However, as our focus is on assessing potential competition impacts in relation to the role of the hyperscalers' as distributors of ISVs' services, we do not consider smaller cloud provider marketplaces further.

expansion for ISVs and weaken the ability of ISVs to compete with the hyperscalers in offering PaaS and SaaS services.

7.27 In principle, there are several routes to market for ISVs in addition to selling through the hyperscalers' marketplaces. ISVs can sell directly to customers (for example, through their own websites) or they can also offer their services via intermediaries.

7.28 Whether ISVs rely on the hyperscalers as a route to market will depend on the extent to which: i) customers view different discovery and purchase routes as substitutes, and ii) hyperscalers control access to particular customer groups. ISVs are less likely to rely on the hyperscalers as a route to market if most customers are using, or are willing to use, alternative purchase channels. However, it cannot be ruled out that ISVs may rely on hyperscalers as a route to market, if there were customers which only used a single purchase channel controlled by hyperscalers and if these customers were sufficiently important for ISVs.

## At present, ISVs do not significantly rely on marketplaces to access customers

7.29 At present, only a small proportion of customers purchase services through marketplaces. The market research found that only 13% of IaaS/PaaS users purchase third-party services through marketplaces.[993] Most marketplace users (75%) use marketplaces for other reasons (such as researching/discovering new services, billing for existing services, or buying first-party services) and do not use marketplaces to buy third-party services.[994] This implies that most customers are purchasing ISVs' services through other channels.

7.30 Consistent with this, we also estimate that only a small proportion of ISVs' sales are made through marketplaces. We estimate that less than 10% of total ISVs' sales (PaaS and SaaS) in the UK are transacted via a marketplace.[995] Even for those ISVs which we understand make a higher proportion of their sales through marketplaces, these sales still make up only a minority of their total sales. This implies that ISVs can and do sell their services via other channels. [✂] noted that "many sellers don't yet view marketplaces as a strategic new customer acquisition channel".[996] [✂] said that while third party services are widely used on public cloud infrastructures a relatively small proportion are currently purchased through marketplaces.[997]

7.31 However, as discussed in Section 4, there are some indicators that marketplaces may grow in importance over time. In addition to the evidence set out in Section 4, we note that SDIA said that, while it generally agrees with our analysis, over time ISVs are likely to become increasingly reliant on the marketplace for customer acquisition, and once that dependency is established, it is probable that hyperscalers will charge fees similar to Apple's App Store or

[993] Context Consulting research data tables, Q45 and Q46. We do not know what proportion of ISV's purchases customers make through marketplaces.
[994] Context Consulting research data tables, Q45 and Q46.
[995] We calculated two estimates of total PaaS and SaaS market size excluding the hyperscalers. One estimate was based on responses to our statutory information requests, IDC PaaS and IDC SaaS data. The other estimate was based on responses to our statutory information requests, Synergy PaaS and IDC SaaS data. We then divided total third-party sales through AWS, Microsoft and Google marketplaces by these estimates.
[996] See [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[997] [✂].

Google's Play Store.[998] Also, [✂] said that, while it may be true that marketplaces are not yet a major revenue source for the hyperscalers, it encouraged us to refer this issue for further investigation by the CMA before AWS's marketplace becomes another tool to maintain its dominance.[999]

7.32 In principle, marketplaces may be more important sales routes for smaller ISVs which lack the necessary sales infrastructure to be able to reach a wide base of potential customers. However, we did not receive evidence on the importance of marketplaces for smaller ISVs.

7.33 Customers differ in how often and how they use marketplaces. The ability to include marketplace spending towards committed spend attracts many customers to use marketplaces where they can.[1000] Other important reasons for using marketplaces include the ability to secure additional discounts, for simplified billing and terms, and uniformity in the management of the product lifecycle.[1001] It has been claimed that "more and more customers just default to buying all their third-party software for their AWS environments through AWS Marketplace", rather than going directly to ISVs.[1002] But others have told us that they do not use marketplaces and that they prioritise direct procurement relationships with vendors.[1003] Other customers make very limited use of marketplaces – or do not use them at all.[1004] Many of these customers prefer to procure directly with vendors in order to control costs. Even those customers that currently make use of marketplaces told us that they are able to use multiple purchase channels, going direct to vendors if they want to.[1005] We have not seen any specific evidence that any customer groups are particularly reliant on marketplaces as a purchase channel.[1006]

7.34 Marketplaces may have an important role in allowing customers to discover new services, allowing ISVs to be exposed to a large number of potential customers, even if ultimately customers purchase those services through a different channel. But the evidence on

---

[998] SDIA response to the interim report, question 5.3.
[999] [✂] confidential response to the interim report, page [✂].
[1000] Including [✂]. See [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]. According to a survey conducted by Tackle, nearly half of all customers consider the ability to draw down on committed spend as the most important benefit of cloud marketplaces, the largest single reason, and 66% of sellers initially listed on marketplaces to tap into buyers pre-existing spend commitments. Tackle 2022, State of Cloud Marketplaces 2022 | Tackle [accessed 7 August 2023].
[1001] See [✂] response dated [✂] to our customer questionnaire, question [✂]; [✂] response dated [✂] to our customer questionnaire, question [✂]; and [✂] response dated [✂] to our customer questionnaire, question [✂]. "Billing for existing services" was the largest single purpose for using marketplaces (64% of marketplace users). Context Consulting research data tables, Q46.
[1002] Protocol, 2022. AWS has a clear advantage among cloud enterprise marketplaces: It has the most customers [accessed 7 August 2023].
[1003] For example, [✂] told us that it would only choose to purchase products through a cloud marketplace if a direct relationship is not possible. [✂] response dated [✂] to our customer questionnaire, question [✂].
[1004] This includes [✂]. See [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[1005] For example, [✂] told us that it is normally able to purchase directly from the vendor. See [✂] response dated [✂] to our customer questionnaire.
[1006] [✂] told us that it has not observed a particular pattern of types of customers procuring through particular distribution channels. [✂] response dated [✂] to s.174 notice dated [✂], question [✂].

whether marketplaces have an important role in discoverability is mixed. The market research suggested that more IaaS/PaaS users use marketplaces to discover new services (25%) than to purchase ISVs' services (13%).[1007] The market research indicated that those which had switched IaaS or PaaS provider completely, and those who had recently started using PaaS services, were more likely to use marketplaces for discovery.[1008] This suggests that marketplaces could have a greater role in discovering service options, which may then be purchased either through the marketplace or elsewhere.

7.35 Customers differ in whether they use marketplaces for discovering new services. [✂] said it uses marketplaces to discover services, but may then buy direct. In contrast, [✂] stated that they use marketplaces to purchase services they are already aware of.[1009] Further, AWS stated that marketplaces are rarely used for casual shopping; and Microsoft said that sophisticated customers are using the marketplace to fulfil specific cloud services needs with solutions that customers have already determined before visiting the marketplace.[1010]

7.36 This implies that some customers are less likely to be using marketplaces to discover new services. The market research finds that just under half of marketplace users who buy third-party services through marketplaces are not currently discovering services through the marketplace – implying that these customers already knew about the service they wished to purchase, and specifically decided to purchase that service via the marketplace route rather than an alternative.[1011] Overall, while some customers use marketplaces as a discovery tool, a material fraction of customers do not seem to rely on marketplaces for discovery.

## We have seen little evidence of harmful behaviour or practices from marketplace operators

7.37 As noted, in general, ISVs do not currently rely on marketplaces to access customers. As a result, marketplace operators are unlikely to have the ability to exploit ISVs seeking to list on their marketplaces. Consistent with this, we have seen little evidence of marketplace owners undertaking any exploitative or exclusionary behaviour in practice.

7.38 Marketplace owners do not require ISVs to exclusively list on their marketplace.[1012] We note, however, that AWS does not allow ISVs to include software or metadata that redirects users to other cloud platforms, additional products, upsell services or free trial offers which are not available on AWS Marketplace for their SaaS-based products. It also prevents sellers from including advertising, promoting or links to purchase services which are not listed in AWS Marketplace. However, this does not explicitly prevent vendors from selling their services through other routes.[1013]

---

[1007] Context Consulting research data tables, Q45 and Q46.
[1008] Context Consulting research data tables, Q46.
[1009] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[1010] AWS response dated 31 October 2022 to s.174 request dated 24 October 2022, page 3; and Microsoft response dated 18 November 2022 to s.174 notice of 21 October 2022, question 24d.
[1011] Analysis of Context Consulting research data tables, Q46. The market research does not tell us how many customers discover through marketplace but then buy direct.
[1012] [✂] told us that listing on a particular marketplace does not prevent it from selling outside of such marketplace. See [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[1013] AWS website. SaaS product guidelines - AWS Marketplace [accessed 7 August 2023].

7.39 Hyperscalers told us that their marketplace mechanisms for organising and displaying solutions to customers do not preference their own services – instead, basing this on search experience, customer popularity or benefits programs.[1014]

7.40 AWS said that it is transparent and neutral with accurate and unbiased product information,[1015] and we note that almost all ([✂]%) of its marketplace sales are third-party cloud services – which suggests it does not self-preference.[1016] Google stated that it [✂].[1017] Google also said that: [✂].[1018] We understand that marketplace owners allow customers to draw down spend commitments when purchasing third-party services via marketplaces – although, in some cases, the overall drawdown through marketplace sales is limited (e.g. [✂]), or is less than the drawdown for purchases of first-party (e.g. [✂]).[1019]

7.41 At present, we have heard few specific concerns from ISVs about the prominence or ranking of their services on marketplaces. However, [✂] said that ISVs' services appear less discoverable than their first-party counterparts. For example, hyperscalers' first-party counterparts feature prominently in their consoles,[1020] whereas ISVs' services need to be discovered and installed from a crowded marketplace. Hence, customers may be less likely to take up ISVs' services.[1021]

7.42 We have heard concerns that marketplace owners may gain access to information which they can use to improve their own competing services, or to develop new services.[1022] In theory, it may be possible for cloud providers to gather information on how successful these third-party offerings are and to then develop similar offerings themselves – either from their role as marketplace owners, or more generally through their role as suppliers of cloud infrastructure to ISVs, which could ultimately raise a barrier to entry to ISVs. But we also understand that each of the hyperscalers has internal access policies which prohibit or

---

[1014] See [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]. Microsoft organises some apps as "Featured Apps" based on [✂] and badges certain offerings as "preferred solutions" if they meet certain criteria. See [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and Microsoft website. Learn about the Microsoft preferred solution badge and Microsoft website. Microsoft Azure Marketplace [accessed 7 August 2023].

[1015] See AWS response dated 28 November 2022 to the s.174 notice dated 24 October 2022, Annex Q33.1.15, page 8.

[1016] See AWS response dated 9 December 2022 to s.174 request dated 24 October 2022.

[1017] See Google response dated 23 November 2022 to s.174 request dated 26 October 2022, annex 33.22, p.1.

[1018] [✂].

[1019] See [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]. We also note that [✂] has raised concerns that differentiated treatment in relation to marketplace spend commitments impacts on its ability to compete (Ofcom / [✂] meeting, [✂]).

[1020] A console is a web application that allows users to access the cloud provider's cloud services. The console can be considered the backbone or basic web infrastructure through which a cloud provider's services can be accessed.

[1021] Ofcom / [✂] meeting, [✂]; Ofcom / [✂] meeting, [✂]. Ofcom / [✂] meeting, [✂].

[1022] [✂] raised concerns that hyperscalers have access to data, such as what products are popular, what customers are using them for, and how much they are charged. It also raised a case where [✂] was able to tell that some of its customers [✂]. Ofcom / [✂] meeting, [✂]. A similar concern was also cited by ACM in relation to Elastic's Elasticsearch and AWS's Elasticsearch Service. ACM, 2022. Market study cloud services (acm.nl), page 63 [accessed 7 August 2023].

restrict service product and engineering teams from accessing transactional data from its marketplace.[1023] We have not seen further evidence on this.

7.43 In fact, we see behaviour suggesting that marketplace owners are seeking to attract ISVs and customers to use their marketplace – benefiting both in the short-term. Both Google and Microsoft have recently reduced their listing fees significantly from as much as 20% to 3%[1024] and we understand that, [✂], the average listing fee AWS receives [✂][1025] and can be low-single digits for some ISVs.[1026] [✂].[1027] As set out in Section 4, marketplace owners often offer incentives for ISVs to sell via their marketplace. Further, we have seen evidence of marketplace owners adjusting their marketplace offerings to respond to ISVs' demands. For example, we understand AWS introduced Private Offers in 2017 (which allow customers to negotiate custom contracts directly with sellers), in order to discourage customers from discovering a service on the marketplace, but then purchasing directly from vendors outside of their marketplace.[1028] Additionally, [✂].[1029] Evidence we have gathered suggests that marketplace owners are considering and responding to changes made by other hyperscalers to their marketplaces and monitoring rival marketplaces.[1030]

## Conclusion on hyperscalers as distributors of ISVs' services

7.44 The evidence that we received during this market study suggests that, at present, ISVs generally do not rely on the hyperscalers to distribute their services to customers. Further, we have seen little evidence that the hyperscalers are engaging in practices within their marketplaces which could act as a barrier to entry or expansion for ISVs.

7.45 That said, the hyperscalers appear focused on increasing the use of their marketplaces by their customers and ISVs (see Section 4, paragraph 4.89), and we anticipate that marketplaces are likely to grow in importance as a route to market for ISVs. We are aware that the balance of the hyperscalers' incentives may change going forward and the risk of foreclosure via marketplaces may increase which is a concern some stakeholders have raised. We consider this risk further in Section 8. In addition, our assessment is based on input that we received from a small number of mainly larger ISVs – and which may not reflect the experience of smaller ISVs.

---

[1023] [✂]. See [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to our follow-up email dated [✂] concerning the s.174 notice dated [✂], question [✂].
[1024] CRN, 2021. Partners Cheer 'Substantial' Microsoft Marketplace Fee Cuts [accessed 27 September 2023] and Google Cloud, 2022. Google Cloud doubles-down on ecosystem in 2022 to meet customer demand [accessed 27 September 2023]. [✂]. See [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[1025] [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; and Ofcom analysis of [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[1026] Protocol, 2022. AWS has a clear advantage among cloud enterprise marketplaces: It has the most customers. [accessed 7 August 2023].
[1027] For example, [✂] has negotiated fees of [✂]%. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[1028] Protocol, 2022. AWS has a clear advantage among cloud enterprise marketplaces: It has the most customers. [accessed 7 August 2023]. [✂] ([✂] response dated [✂] to the s.174 notice dated [✂], [✂]).
[1029] [✂] response dated [✂] to the s.174 notice dated [✂]; [✂].
[1030] For example, see [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂]; [✂].

# 8. Competition concerns we have identified in this market study

8.1 In this market study, we have considered the extent to which a matter in relation to the provision of cloud infrastructure services in the UK has or may have effects adverse to consumers. This could lead to higher prices, lower quality products or less innovation.

8.2 Cloud computing is an increasingly important input to the different elements that make up the internet and how online services are developed and delivered to UK consumers. This includes every type of consumer activity which takes place online, covering services such as social media, streaming, and communications services.

8.3 Cloud services are also an input in sectors that produce products and services that are not (entirely) digital. This is the case particularly for cloud products that can be relevant to any type of business, such as SaaS products for employee and customer management or PaaS products which can power AI/ML solutions.

8.4 Cloud is also a cornerstone of recent technological innovations. From data science to AI, many of the cutting-edge developments in the way software is transforming how we live our lives, run our businesses, and engage with our public services, is operating from the cloud.

8.5 This means that competition concerns in cloud services can have wide-ranging effects across both online and other products and services that UK consumers buy. Such effects could take different forms:

a) Where business customers face higher costs to source the cloud services they need, they may pass these on in ways that will ultimately lead to higher prices for the products and services UK consumers buy.
b) Limits on entry and expansion and innovation can directly affect the quality and range of choices UK consumers have of online services powered by the cloud.

8.6 We have carried out our assessment by looking in detail at how competition works in cloud infrastructure services (Section 4) and then examining the various barriers to effective competition that may exist (Sections 5 and 6). Based on our findings, we set out our views on how well competition is currently working in cloud infrastructure services. Within this section we also discuss the potential implications of a lack of competition on UK businesses and other organisations.

## There are factors which inhibit effective competition in cloud infrastructure services

### Cloud providers predominantly compete to attract customers into their ecosystems

8.7 Cloud providers have developed ‘ecosystems’ which combine first- and third-party products across IaaS and PaaS. These cloud ecosystems benefit customers by allowing them to source a wide range of products from one place and to easily combine them to build their IT

solutions. An important dynamic is that providers compete at the ecosystem level, where they vie to supply all or most of the cloud needs of customers.

8.8 As detailed in previous sections, once a customer is established with a cloud provider, there are material barriers to switching and multi-cloud. As a result, much of competition in cloud infrastructure is currently centred around acquiring customers when they first migrate into the cloud. This is evident in cloud providers' strategies, where they use discounts and other incentives to draw customers into their ecosystems. Providers then have the potential to earn higher revenues later as the credits and free trial periods come to an end and acquired customers begin to take up further services.

8.9 Our evidence indicates that AWS's and Microsoft's ecosystems are viewed by customers as serving the widest breadth of use-cases and access to different functionalities. They are also likely to benefit most from economies of scale and network effects due to their large established customer bases. This is reflected in the fact that AWS and Microsoft accounted for [✂]% [70% to 80%] of UK revenues across IaaS and PaaS in 2022.[1031]

## There is limited scope for workload competition amongst cloud providers

8.10 In the interim report, we noted that there is some competition for narrower components of customers' cloud needs, which we refer to as 'workload competition'. In their responses, Microsoft[1032] and AWS[1033] stated that there is a considerable amount of workload competition as evidenced by the number of customers with multi-cloud deployments, and it is not the case that once a customer has cloud spend with one provider, it will not award new workloads to another.

8.11 We agree that there is scope for some competition between cloud providers for new workloads, either because customers migrate additional workloads to the cloud or because their needs evolve. We have seen examples of some larger customers that already have a primary cloud provider adding a secondary provider to serve a new use-case. However, as discussed in Section 4, our evidence suggests that workload competition is mainly limited to more siloed use-cases (where applications/data do not have to be integrated, or only to a limited extent, with those of the primary provider) and/or to back-up certain workloads. We found little evidence of customers today using more integrated types of multi-cloud and we consider this in part reflects the barriers that we identify in Section 5.

8.12 While workload competition for more siloed use-cases is likely to have some benefits (e.g. customers have IT staff with skills to work on multiple clouds which may increase their ability to switch), there are limitations to the level of competitive constraints this offers. This is because siloed multi-cloud is only likely to be feasible for a limited number of use-cases and is not suitable where customer IT architectures require some amount of integration. Consistent with this, our evidence indicates that where customers have multi-cloud architectures, their spend is generally concentrated around a primary provider with only a small number of niche services being taken from a secondary provider (as discussed in Section 6). Moreover, due to the barriers to switching and multi-cloud we identified, there may be limited competitive pressure for siloed workloads once built out on a cloud.

---

[1031] See Section 4 for further detail on market shares.
[1032] Microsoft response to the interim report, page 37, paragraph 139.
[1033] AWS response to the interim report, page 5, paragraphs 10-11.

8.13 We also see some workload competition within cloud providers' ecosystems, as customers can choose between first-party services or ISV services for some PaaS types. As set out in Annex 1, we estimate that ISVs collectively account for [✂]% [30% to 40%] of UK PaaS revenues in 2022. However, as discussed in Section 6, our analysis suggests that no single ISV has a share greater than 5% in PaaS. We also think there are reasons why the broader competitive constraint that ISVs can exert on the ecosystems of the cloud providers is limited.

8.14 First, due to the convenience of purchasing cloud services in bundles, cloud providers may have some advantages over ISVs. As noted in Section 6, ISVs tend to offer products in one or a small number of PaaS product categories, whereas cloud providers are present across all PaaS categories. In addition, we have identified some technical and commercial factors that may disadvantage ISVs when competing with first-party products (as discussed in Sections 5 and 7).[1034] The fact that hyperscalers account for the majority of PaaS sales is consistent with many customers opting to source most of their cloud needs using first-party products from a cloud provider.

8.15 Second, ISVs typically do not have their own physical infrastructure, and therefore do not provide IaaS. As a result, competition from ISVs will not constrain the IaaS elements within customers' workloads. Even where customers choose an ISV for PaaS services, they will typically need to use them in conjunction with IaaS elements (e.g. compute and storage),[1035] which likely represent a significant share of the value of many workloads. As set out in Section 3, in 2022 UK IaaS revenues were £[✂] [£4.0 to £4.5] billion as compared to £[✂] [£2.5 to £3.0] billion PaaS revenues. The level of constraints on IaaS services depends on competition between cloud providers, which we discuss below.

8.16 We therefore maintain our view that there is competition for narrower components of cloud ecosystems – e.g. between cloud providers for new workloads across IaaS and PaaS, and within ecosystems between cloud providers and ISVs for PaaS – but that at present the scope of this workload competition is limited.

## While customers benefit from cloud infrastructure services, we see evidence that competition is not working well

8.17 We can see some positive outcomes for customers that are likely to be driven by the competitive dynamics in cloud infrastructure services. These include:

a) **Innovation**: Quality, number of features and the range of products are key parameters customers consider when choosing a cloud provider. In response, cloud providers are continuously innovating to attract new customers by differentiating their offerings, and increase the usage of existing customers. Cloud providers innovate to expand the features of existing products and develop new ones, alongside investments in making physical infrastructure components (e.g. chips) more powerful and efficient. Cloud providers are also at the forefront of developing cutting-edge technologies such as AI. As

[1034] Factors that may disadvantage ISVs may include: (i) the need for customers to manage separate ancillaries to make full -use of ISV services, this increases customers' hassle costs; (ii) cloud providers may not make new functionalities and innovations that it introduces for its first-party services immediately available to ISV services; and (iii) hyperscalers' first-party services can be found easily by customers in their consoles, whereas ISV services need to be found separately within marketplaces.

[1035] Depending on the service, compute/storage elements can be packaged within ISV PaaS services or customers can purchase these directly from the cloud infrastructure provider.

set out in Section 4, in response to our interim report, customers have highlighted the benefits they get from accessing new services as providers continue to improve their offering, and cloud providers have told us that they pass on cost savings to customers.

b) **Customer choice of ISVs**: Customers have a broad range of ISV products on different clouds to choose from. The hyperscalers enable some ISVs by providing them with technical support to integrate their services and engaging in shared selling activities. We think this is because uptake of ISVs benefits the hyperscalers through greater use of their underlying infrastructure by existing customers. They may also want to expand the range of functionalities available to attract new customers.

c) **Customer choice of open-source products**: The hyperscalers have responded to customer demand for open-source technologies, at least in part. For example, each of the hyperscalers developed its own container services and container orchestration platforms, including some based on open-source products (e.g. Kubernetes). Other examples include hyperscalers using Linux OS and offering some database services based on open-source technologies such as PostgreSQL, MariaDB/MySQL and RedisCloud. In addition to providing customers with more choice, the use of common open-source technologies across hyperscalers can lower barriers to switching for customers by providing a certain level of standardisation.

d) **Pricing trends**: As detailed in Section 4, data submitted by AWS and Microsoft suggests that list prices paid by PAYG customers for their core cloud infrastructure services have either remained stable or decreased in recent years. The data also suggests that average net prices for these services (which incorporate all types of discounts) have fallen faster than list prices over the same timeframe, although the list price and net price series are not always perfectly comparable. As explained in Section 4, AWS told us that it has passed on global cost savings to customers through price reductions.

8.18 Despite providers competing to win new customers and some more siloed workloads, our study has identified various indications that the market is not working well. This appears to be driven by some customers facing significant barriers to switching or multi-cloud once they become established on a particular cloud. In particular, we have found that:

a) **Switching levels are low:** our market research found that 'difficulty and expense of switching provider' was respondents' biggest single concern about the cloud infrastructure market (59% of respondents).[1036] While our market research found that switching is less cumbersome for some, we found that only c.20% of customers have switched providers[1037]– which is likely to include switching within a cloud (e.g. between first-party and third-party services), or switching between on-premises IT/private cloud and public cloud. Our large customers questionnaire was consistent with this, indicating that it is uncommon for customers to switch away from a provider completely due to the time and cost required. Following the interim report, [✂] submitted quantitative analysis based on its UK cloud customer usage and revenue data which it considers shows that customers can and do switch provider. However, we note that the churn rates calculated by [✂] are consistent with the results of our market research (see Section 4 and Annex 3 for further details).

b) **Integrated multi-cloud is uncommon:** Our evidence indicates that it is rare for customers to use multi-cloud architectures where applications are closely integrated

[1036] Context Consulting research report, slide 131.
[1037] Context Consulting research data tables, Q47.

and significant volumes of data are being transferred between them. Our market research found that at least half of respondents used a single cloud provider.[1038] Our engagement with large customers and intermediaries indicated that, of the customers using multi-cloud, this was for use cases that are at the more siloed end of the multi-clouding spectrum. Since the interim report, [✂] has submitted analysis on this issue. However, we explain in Section 4 and Annex 3 why the approach used by [✂] does not provide a meaningful measure of the prevalence or extent to which its customers multi-cloud. We therefore maintain our view that integrated multi-cloud is at present uncommon.

c) **AWS and Microsoft have consistently high shares of supply:** As set out in Section 4, between 2019 and 2022, we estimate that AWS's and Microsoft's combined share of UK IaaS and PaaS revenues increased, and in 2022 stood at [✂]% [70 to 80%]. Over the same period, Google grew steadily by [✂] each year, reaching [✂]% [5 to 10%] in 2022, but its share remains considerably smaller than the market leaders. Beyond the hyperscalers, the next largest providers have considerably smaller shares. Specifically, we note Oracle's share of UK IaaS and PaaS revenues has not grown significantly since it started offering cloud services in 2016, and its UK share remains around [✂]% [0% to 5%]. Looking at IaaS only, we estimate that in 2022 AWS and Microsoft represented approximately [✂]% [80% to 90%] of UK IaaS revenues, a share that has increased by around [✂] [5 to 10] percentage points since 2019.

d) **High profitability of AWS and Microsoft Azure:** Returns in excess of the cost of capital over a sustained period could be an indication of limitations in the competitive process.[1039] As set out in Section 4 and Annex 2, AWS's returns have been significantly above its cost of capital in every year since 2014; and between 2018 and 2022 its return on capital employed averaged 40%. Microsoft Azure's returns have exceeded its cost of capital in each of the last three years and are increasing. We recognise that limitations in the competitive process are not the only possible explanation for high returns, as other factors, such as the reward for innovation and efficiency, could also contribute. However, as explained in Annex 2, we do not consider it likely that these other factors can fully explain the persistence and magnitude of returns above WACC, especially for AWS.

8.19 In combination with the barriers we have identified in Section 5, high profitability, a concentrated market structure and limited levels of switching and integrated multi-cloud are likely indicators that cloud providers are not facing sufficiently strong competitive constraints.

[1038] As set out in Context Consulting research report, slide 75, 52% of IaaS/PaaS users reported using more than one IaaS/PaaS provider. This is likely to overstate the use of more than one public cloud provider. This is because, in addition to using multiple public clouds, some respondents who use more than one IaaS/PaaS provider may be combining: (i) the products of an ISV and public cloud provider on the same cloud; (ii) private and public cloud solutions (i.e. hybrid cloud); or (iii) two private cloud providers. We discuss this further in Annex 3.

[1039] Competition Commission, 2013. Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraph 118 [accessed 6 September 2023]

## Our assessment of barriers to switching and multi-cloud contribute to our view that the market is not working well

8.20 In light of indicators that the market is not working well, we have examined the potential reasons why customers find it difficult to switch or multi-cloud.

8.21 First, we have found that there are inherent technical barriers to switching cloud infrastructure services. This is common to many IT markets – it stems from technical differences in the way providers have designed their cloud stacks. Such technical differences may be the result, in part, of cloud providers competing for customers by differentiating their services, e.g. as providers innovate to introduce new products and design their systems in ways they consider to be technically efficient. However, these differences result in customers facing a low degree of interoperability and application portability across cloud providers as they use different proprietary cloud technologies (e.g. APIs, protocols, workflows, programming languages and data formats). Consequently, customers wishing to switch will face costs from needing to re-engineer their applications so that their applications can 'talk to' and run in the target cloud.

8.22 Similarly, we have found that technical barriers can also raise the cost of customers deploying applications across multiple clouds. Customers that wish to use an integrated multi-cloud approach may face particular challenges and costs. For example, additional technical effort, such as deploying adaptors, may be needed to allow cloud solutions which deploy different clouds in an integrated way. In addition, the differences in cloud technologies across providers requires customers to develop specific skills for each cloud. As a result, a customer wanting to multi-cloud or switch to a new provider will need to invest time and resources into retraining or hiring new staff to be able to operate in that new cloud.

8.23 Second, we have found that there are certain features of the market that may further raise barriers to switching and multi-cloud:

a) **Egress fees:** Egress fees are likely to be a significant barrier to customers using integrated multi-cloud, particularly where this requires large volumes of data to be transferred between clouds. Egress fees also act as a barrier to switching. As well as the one-off costs of transferring data to the gaining cloud provider, customers are likely to host services on multiple clouds during the switching period which can require data to be transferred between them. The evidence we have gathered indicates that AWS, Microsoft, Google and IBM are currently setting egress fees at a level above what is needed to recover costs.

b) **Limiting interoperability/portability of their services:** Some technical barriers to switching and multi-cloud may not be justified: (i) ancillary services (such as security, access management, monitoring, billing) are more differentiated than they need to be; (ii) first-party cloud services are not always fully functional when used in combination with third-party cloud services; (iii) solutions that can mitigate data latency when connecting clouds are not widely available across clouds; and (iv) cloud providers may not be fully transparent about the compatibility of their cloud infrastructure services with third-party services.

c) **Committed spend discounts:** Cloud providers negotiate committed spend discounts with their largest customers, where the percentage discount a customer receives increases as the amount they commit increases. While such discounts can benefit customers through lower prices, the way in which the discounts are structured

encourages customers to purchase all or most of their cloud needs from a single provider. This discourages the use of multiple providers for customers' existing and incremental workloads.

8.24 Each of these features can individually make it less attractive for customers to switch and operate integrated multi-cloud architectures. For example, customers have told us that egress fees are the biggest single challenge to using multi-cloud (45% of respondents)[1040] and there is evidence of customers seeking alternative cloud architectures to minimise or avoid egress fees. We have also heard from large customers that committed spend discounts are an important commercial consideration for them. We have seen evidence that the threat of losing a discount over their existing spend can serve as a strong incentive for customers to bring new spend under their commitment with their existing provider.

8.25 There is also the potential for these features to work in combination and have a greater impact on customers' decisions. To illustrate, customers may be discouraged from implementing multi-cloud architectures due to, for example, the total financial cost of having to pay egress fees to transfer data between clouds and losing a committed spend discount where splitting workloads between providers causes them to miss out on their spending commitments. These incentives can ultimately result in 'path dependence', where once a customer makes the initial choice of cloud provider, they are more likely to deploy future workloads from the ecosystem of that same provider.

8.26 Furthermore, any barriers to adopting integrated multi-cloud can also make it more difficult for customers to fully switch cloud provider. This is because many customers who wish to switch providers cannot switch all workloads between providers instantaneously. Rather, they must migrate workloads gradually from one cloud to another over a significant period of time (e.g. several months) and run workloads/applications in an integrated way across different clouds during this period. As a result, any features that make it more difficult for customers to adopt more integrated multi-cloud architectures will also likely make it more difficult for them to switch.

## A material share of customers may face barriers to multi-cloud and switching

8.27 The extent to which customers are affected by the barriers we have identified will depend on their individual needs. Some customers may be able to switch relatively easily as they take few products that are more easily ported between cloud environments (e.g. basic IaaS products). Customers may also be able to reduce technical barriers to switching/multi-cloud to some extent by building their solutions using cloud-neutral design principles. For example, by using container services or open-source services that are not specific to a particular cloud environment. In both cases, this is only likely to be feasible for customers with few applications and simple needs, such as smaller start-ups.[1041]

8.28 However, our evidence suggests that a large portion of the market is likely to have more complex needs and is likely to face high barriers to switching or adopting more integrated multi-cloud architectures once they have chosen their primary provider. This is relevant to

[1040] Context Consulting research data tables, Q31; and Context Consulting research report, slide 79.
[1041] Where customers operate a larger number and variety of applications, or across multiple teams within their organisation, building cloud-neutral architectures is often impractical, due to the additional complexity, general time constraints and the amount of centralised coordination needed.

both PAYG customers and those that negotiate contracts. This is consistent with a range of evidence collected during the study:

a) Our market research found that 'difficulty and expense of switching providers' was the top concern of PAYG and contract customers about the cloud infrastructure market. It was cited by more than half of respondents (59%), with concerns about egress fees and interoperability also ranking highly (55% and 52% respectively).[1042]
b) RFI data we collected from the hyperscalers shows that customers spending more than £10k per year (representing [✂]% of revenues) consume at least [✂] [10-20] first-party proprietary products and at least [✂] [5-15] PaaS services.[1043] Based on our technical analysis, we consider that customers taking this number of products across IaaS and PaaS would be more likely to face a high degree of technical complexity if they wanted to switch or multi-cloud. This is consistent with our market research, which found that 84% of respondents have between 4 and 10 use-cases and were more likely to cite technical challenges as a barrier to switching compared to those with fewer use-cases.[1044]
c) Evidence from one hyperscaler indicates that only a very small number of customers are able to negotiate egress fee discounts in advance (less than [✂] have negotiated a discount). The remaining [✂]% of customers pay the list prices for egress or do not exceed their free tier allowance (but would pay the list price were they to exceed this). These customers account for over [✂]% of cloud revenue [✂].[1045]
d) Hyperscaler data indicates that committed spend discounts are relevant to a large proportion of the market, accounting for [✂] ([✂])% of AWS and Microsoft's UK cloud revenues.[1046]

8.29 In relation to some of these features, more mature and larger customers may be particularly affected. Our market research indicates that technical barriers may be more significant for more mature cloud users.[1047] This may be because these customers are more likely to have large numbers of applications and/or use various proprietary services offered by their cloud providers, which add to the complexity of switching cloud provider. Evidence gathered from the hyperscalers is consistent with this, indicating that customers spending more than $1m a year (accounting for [✂]% of hyperscaler revenue) on average take at least [✂] [30-40] first-party proprietary services and more than [✂] [20-30] PaaS services.[1048] Similarly, committed spend discounts are more relevant for larger customers, as these pricing arrangements are widespread among hyperscaler customers with more than $[✂] spend a

[1042] Context Consulting research data tables, Q63; and Context Consulting research report, slide 131.
[1043] See Annex 4 for further details.
[1044] See: Context Consulting research report, slide 37; and analysis of Context Consulting research data, Q52. On average 47% of customers with 0-3 use-cases selected at least one of these technical challenges as a barrier to switching: data portability, application portability and interoperability. This percentage rises to 60% for customers with 4-10 use-cases. This compares the percentage of customers with 0-3 use-cases with the average across customers with 4-5, 6-7 and 8-10 use-cases. The difference is even starker after excluding customers using private cloud only. In this case, 47% of customers with 0-3 use-cases selected at least one of these technical challenges as a barrier to switching: data portability, application portability, interoperability. This compares to an average of 61% for customers with 4-10 use-cases.
[1045] See Section 5 "Data egress fees" for further details.
[1046] See Section 5 "Committed spend discounts" for further details.
[1047] For example, 58% of companies established for more than 2 years indicated at least one technical barrier to switching vs 49% of companies established for less than 2 years. Analysis of Context Consulting research data, Q52.
[1048] See Annex 4 for further details.

year, and cover the vast majority of hyperscaler customers with annual spend above $[✂].[1049] While some of these larger customers may be able to achieve material discounts on egress fees, our evidence suggests this is a minority of customers who nonetheless still face other material barriers.

# The barriers we have identified can harm cloud customers, and ultimately consumers

## Customers that face barriers to switching and multi-cloud can be harmed

8.30 In the interim report, we highlighted concerns that customers who face significant barriers to switching and multi-cloud have a limited ability to credibly threaten to switch all or some of their-workloads to new providers. This can reduce the competitive pressure on their existing provider, such that they can exercise a degree of market power in respect of the existing and incremental workloads of a material share of existing customers. This could lead to higher prices compared to what would be the case if customers could switch or multi-cloud more easily. Customers may also be harmed if there is a better product on offer by a competitor (for example, a more innovative solution) and they cannot take-up these services due to these barriers.

8.31 In response, AWS suggested that there is competition for new workloads,[1050] and Microsoft noted that competition between suppliers for the next cloud workload means that they have no incentive to increase prices for existing customers as this would risk losing future business.[1051] While we agree that there may be greater scope for competition where the new workloads of customers are sufficiently siloed, we have found that competition for new workloads is dampened because of barriers to switching and multi-cloud. As such, we consider the threat of moving new workloads to be less of a disciplining effect than has been suggested by the market leaders. This can result in potential harms to both customers who negotiate contracts with the hyperscalers as well as PAYG customers.

### Harms to customers who negotiate prices and PAYG customers

8.32 We remain of the view that barriers to switching and multi-cloud can make it difficult for customers to mix and match products across cloud providers, which can lead them to settle for lower quality alternatives from their existing provider rather than best in breed rival products.

8.33 Moreover, for customers who negotiate contracts with the hyperscalers, barriers to switching and multi-cloud are likely to weaken their bargaining power once they have chosen their provider. The importance of bargaining power in influencing negotiations is acknowledged by [✂] and reflected in the wide variation we observe in discount outcomes for customers negotiating committed spend discounts. This suggests that there is scope for worse discounts for customers who have a weaker bargaining position due to barriers to switching and multi-cloud.

[1049] See Section 5 "Committed spend discounts" for further details.
[1050] AWS response to the interim report, page 5, paragraphs 10-11.
[1051] Microsoft response to the interim report, pages 7 and 8.

8.34 PAYG customers do not typically negotiate individual terms, instead they pay the publicly available list prices and may receive standard discounts off these list prices. The presence of barriers to switching and multi-cloud could dampen competition for these customers, leading to higher prices. If many existing PAYG customers cannot easily switch to a rival that offers a better deal, a hyperscaler need not offer a lower price to retain these existing customers. Similarly, hyperscalers are less incentivised to lower list prices as this may not attract workloads from a rival's customers because of switching costs. Hyperscalers may also be less incentivised to negotiate with PAYG customers on deviations from list prices where these customers have weaker outside options.

8.35 Since new and existing PAYG customers pay the same list price, the presence of new customers could serve as a disciplining mechanism that limits the hyperscalers from exploiting customers who face barriers to switching and multi-cloud. However, there are limitations to the extent of this disciplining effect. When considering whether to reduce list prices, a hyperscaler is considering both the gain of attracting customers new to cloud and the negative financial impact of having to lower prices for its existing customers. The strength of the constraint from new customers depends on several factors, including the relative size of new and existing customer bases.

8.36 Given the material existing customer base of the market leaders, the negative financial impact of lowering prices to existing customers could be material. While cloud is a growing market and there will always be new customers, the number of new customers that could potentially discipline PAYG prices will decrease once more customers become established on the cloud.

8.37 Finally, not all new customers will necessarily place a strong constraint on list prices, for example if they need services that only certain hyperscalers offer. Hyperscalers could also compete for new customers through other means (such as offering one off benefits in the form of cloud credits) and not necessarily through PAYG list prices.

### Customers are unlikely to fully offset the harms from lock-in when they first move into the cloud

8.38 In response to our interim report, some hyperscalers argued that the cloud infrastructure market was characterised by "sophisticated buyers". Even if it were true that customers faced barriers to switching, some hyperscalers argued that they would have costed these in when making their initial design choices.[1052] These submissions reference various industry reports which recommend taking into account vendor lock-in when designing cloud architecture.[1053] Following this logic, any worse outcomes for locked-in customers (which the hyperscalers dispute) would have been foreseen by customers and factored into their initial procurement decision.

8.39 We recognise that some customers can protect themselves to some extent from the negative consequences of the barriers we have identified when they first move into the public cloud. We have heard examples of actions that some customers have taken to try and offset potential harm. However, these actions are costly and time consuming, both in the initial design and ongoing maintenance of their cloud architecture. They can also put

[1052] AWS response to the interim report, page 12, paragraph 27; and Microsoft response to the interim report, page 9 and page 25, paragraph 90.
[1053] See Microsoft response to the interim report, page 25, footnote 92.

artificial limitations on the functionality of the cloud architectures they may build as a result. Therefore, this is not a feasible or attractive solution for many customers.

8.40 In addition, customers face challenges in anticipating their future needs at the point of first moving to the cloud. Customer needs evolve as organisational needs change and as the market evolves and new capabilities are launched, customers will inevitably look to adjust their solutions in response. Customers are therefore unable to build solutions at the outset that meet all new requirements that may emerge or negotiate on prices up front to offset market power on future unknown usage.

8.41 Our engagement with customers indicates there are other factors that could have a bearing on the initial purchase decision. As discussed in Section 5, we have heard from some customers that they struggle to predict their future spend due to a lack of pricing transparency.[1054] We have also seen evidence to suggest that challenges with operating multi-cloud environments and switching are often discovered in practice and not anticipated in advance.[1055] Furthermore, many customers may face bargaining power imbalances with the hyperscalers and be unable to negotiate terms to properly protect themselves against potential future harms.[1056] PAYG customers, who account for [✂] of AWS and Microsoft customers,[1057] are typically unable to negotiate at all.

8.42 For these reasons we do not believe cloud customers can be reasonably expected to fully insure themselves against future harm from barriers to switching or multi-cloud (or harm from 'lock-in') when they first migrate to the cloud.

## Evidence of harm from barriers to switching and multi-cloud

8.43 In our interim report we pointed to a number of factors that suggest barriers to switching and multi-cloud are creating harm for some customers, including: (1) customer feedback on being effectively "locked-in"; (2) customers reporting rising prices in our market research; and (3) profits above WACC for AWS and Microsoft.

8.44 Following the interim report, we engaged with a number of customers who believe that they are in a weaker bargaining position because they cannot credibly switch all or most of their workloads away from their current provider.[1058] We also observe wide variation in the discounts received by committed spend discount customers, agreeing to similar commitments. This suggests that there is scope for worse outcomes for customers who have weaker bargaining power, which could be due to them facing greater barriers to switching and multi-cloud. However, we are not able to test this conclusively with the available data. Further details are provided in Annex 5.

8.45 Barriers to switching and multi-cloud can give providers a degree of market power over existing customers that cannot easily switch or multi-cloud, creating the potential for high

---

[1054] Context Consulting research report, slide 131.
[1055] Context Consulting research report, slide 79. For example, [✂] only discovered the challenges associated with managing both Google and AWS IAM in an integrated fashion several years after their initial cloud migration, when they began the process of adding Google to their architecture. See Section 5 "Technical barriers". We also heard from one customer [✂] that challenges in anticipating future requirements mean that they are unable to see how their initial choices may impact their ability to switch and multi-cloud in the future. Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[1056] See Annex 5 for more detail.
[1057] See Section 4 for more detail.
[1058] [✂] response to the call for inputs, page [✂]; [✂] response dated [✂] to s.174 request dated [✂], question [✂].

profits. As detailed above, we estimate that both AWS and Microsoft Azure earn profits above their WACC. The evidence of persistent and high profitability for AWS, and material increases in Microsoft Azure's profitability more recently, is consistent with them having a degree of market power over some customers and there being limitations on the competitive pressure being placed on hyperscalers.

8.46 In response to our interim report, two hyperscalers, [✂] and [✂], submitted analysis showing that prices typically have not increased upon renegotiation for customers with privately negotiated commitment spend discounts, and therefore [✂].[1059] This differs from the findings from our market research, where a material share of customers who renewed their contracts reported having faced price rises.[1060]

8.47 We acknowledge that the data presented by the hyperscalers does not support our survey finding that price rises are common for customers renegotiating a contract with their provider.[1061] However, the analysis presented by the hyperscalers does not imply that there is no harm to existing customers. We have identified methodological limitations with the analysis, which means we are unable to test whether customers that face barriers to switching and multi-cloud face worse terms than they would have been able to achieve otherwise. Further, we consider that it is unlikely that the datasets provided by the hyperscalers will give us a full account of the extent to which customers are getting better or worse deals upon renewals. This is because the data only includes committed spend discounts and is unlikely to capture all the elements of pricing relevant to committed spend discount customers, such as other forms of discounts they may be able to negotiate on (e.g. service-specific discounts). We set out these limitations in more detail in Annex 5.

## There is a risk that the extent of competition may deteriorate in the future

8.48 Since launching its first cloud service in 2006, AWS has been investing in its product range and network of data centres to take advantage from growing demand, and in later years to fend off competition from rival cloud providers to attract new customers. As the only material provider of cloud services for several years, it has benefited from its position as a first mover to gradually increase its investment while expanding its customer base. As a result, AWS has been able to maintain positive profits while investing to stay ahead as rivals entered the market.

8.49 In contrast, barriers to entry and expansion were more material for those that entered later. This has made it harder for rival cloud providers to attain scale and operate profitably. Unlike AWS, they had to challenge an incumbent who had an established product range and data centre footprint and was already benefiting from economies of scale and network effects. Rival cloud providers also faced greater barriers to acquire customers as switching costs can be significant once a customer has chosen their cloud provider, making it difficult for them to gain scale by competing for the existing and new workloads of customers already on AWS's cloud.

---

[1059] [✂]; [✂].

[1060] Context Consulting research report, slides 101 and 102.

[1061] The price rises captured as part of our research could reflect other changes or issues, and the findings may be more reflective of pricing outcomes for customers without committed spend discounts as only 6% of respondents had a committed minimum spend in place.

8.50 Regardless of these challenges, continued growth in demand has encouraged several cloud providers to invest in building a product range which can serve the needs of many customers, and offer these services across a network of data centres with a global reach. While these cloud providers have been able to attract customers and expand their revenues to some extent, many of them are far behind AWS and have yet to reach a scale where they make material profits. Only Microsoft has managed to grow sufficiently to rival AWS in terms of its established customer base, and to attain a scale at which it is able to earn substantial profits from its cloud infrastructure services. In doing so, Microsoft appears to have benefited from its leadership position in software markets.

## There is a risk that the market becomes concentrated with a few cloud ecosystem only

8.51 The barriers to switching and multi-cloud we have identified can limit scope for workload competition, making it more difficult for cloud providers to compete for components of customer demand. While smaller cloud providers have been able to enter the market, we are concerned that these barriers to workload competition will harm their growth and ability to challenge the market leaders in the future. We consider that there is the potential for the market to become more concentrated around the market leaders, with scope for the harm to be long lasting. While there is no certainty on how the market will evolve, there are several factors which point to a substantial risk of such an outcome materialising.

8.52 AWS and Microsoft have already pulled ahead significantly in terms of their established customer base for cloud infrastructure services. This is likely to offer them cost advantages due to economies of scale, and they benefit most from network effects in terms of attracting ISVs and new customers onto their ecosystems. They may also have other advantages, for example if their large installed customer base provides them with secured revenue streams to invest and improve their quality and scope of offerings. As a result, many new customers are likely to favour the ecosystems of AWS or Microsoft.

8.53 At the same time, we observe that barriers to switching and integrated multi-cloud are sufficiently material for some customers today, leading them to single source or to concentrate most of their usage with the market leaders. This is particularly problematic where costs like egress fees, or unjustified limitations on interoperability, artificially increase the hurdles for smaller providers to compete for customers' workloads. The use of committed spend discounts may further impede the ability for smaller cloud providers to compete effectively for components of a customer's demand.

8.54 This suggests that the ability of smaller cloud providers to gain scale by competing for the business of new customers or the market leaders' existing customers is likely to be limited. This will affect an increasing number of customers in future, as they expand their cloud infrastructure with a primary provider and face barriers to switching their workloads to rival providers. This effect may also be aggravated if difficulties to grow their customer base also makes it harder for smaller cloud providers to attract ISVs into their ecosystems.

8.55 In addition, we anticipate that there will be fewer new customers to compete for once most existing businesses have moved their workloads into the cloud. This means it could become even harder for smaller cloud providers to grow their scale by targeting new customers who have narrower or specialised product needs.

8.56 We therefore disagree with [✂][1062] view that there is no risk of the market concentrating further because smaller providers can use innovative products, or target specific customer needs, as a basis to compete effectively and grow. The barriers to workload competition we have identified risk undermining the scope for smaller cloud providers to gain material scale in this way and close the gap to the market leaders.

8.57 If it becomes more challenging for smaller cloud providers to gain scale and maintain their investment, then this could further weaken their ability to compete effectively with the market leaders going forward. There could be long lasting impacts if this leads the market to become more concentrated towards a few cloud providers, with barriers to switching and multi-cloud allowing the market leaders to entrench their positions and avoid competing vigorously with each other.

8.58 The risk of concentration concern is consistent with the observed shares of supply, which shows that the market has remained concentrated towards the market leaders in recent years. As discussed above, while Google is often positioned as the closest challenger and has been gaining share, it remains far behind in terms of size and the next largest providers have considerably smaller shares.

8.59 In this context we note that AWS has already been able to maintain a return on capital which has been materially above its WACC for several years, and it would appear to have been more able to resist demands for greater interoperability than smaller cloud providers. Microsoft's Azure returns have been trending upwards and are now above our estimate of WACC. Given barriers to entry and expansion, we anticipate that its profits may increase further, and like AWS it manages to continue attracting customers while being less interoperable than smaller rivals.

### Market concentration towards a few ecosystems may further inhibit effective competition

8.60 Ultimately, whether we arrive at a market concentrated around the two market leaders only, or a market with limited competition between a few ecosystems, the barriers to switching and integrated multi-cloud we have identified have the potential to lead to worse outcomes over the longer-term. We consider the forward-looking concerns to be particularly important given the dynamic nature of cloud services, which is characterised by significant scope for ongoing improvement of existing products and the development of services that cater to new use cases. Furthermore, the development of AI is underpinned by cloud computing and AI is also expected to enhance the functionality of software run in the cloud. This means that new workload opportunities will continuously emerge, whether it be from new or existing customers. We believe it is critical that there is lasting competition amongst cloud providers for these emerging opportunities.

8.61 We are concerned that the barriers we have identified will result in cloud providers facing limited incentives to compete for each other's existing customer base. While competition between cloud providers for the new and emerging demand of existing customers could drive competition and innovation, this won't happen if cloud providers are unable to credibly compete for the emerging workloads of existing customers due to barriers to multi-cloud. Since it is difficult for many customers to switch existing workloads to a rival cloud, we cannot rely on switching to spur innovation either. Furthermore, as the market matures

[1062] See [✂].

and there are fewer new customers to compete for, incentives to invest in innovation may be weaker if there are fewer new customers to monetise that investment.

8.62 In addition, should we arrive at a more concentrated market these concerns are magnified as customers may be denied the increased competition, innovation, and product differentiation that additional cloud providers can offer. Weaker competitive constraints on the market leaders from the smaller cloud providers may dampen price competition by limiting the outside options of customers. A more concentrated market may also chill innovation if smaller cloud providers are less able to compete for workloads of new and existing customers, therefore reducing the incentives of market leaders to invest and improve their product offering in response.

8.63 While ISVs play an important role in developing innovative PaaS solutions, they are unable to challenge cloud providers across their portfolio and particularly in relation to innovation in the underlying infrastructure. Competition from ISVs is therefore unlikely to be sufficient to discipline the market leaders if we were to arrive at a more concentrated market. In fact, we explain in the following section how such a market structure may also result in more limited constraints of ISVs at the PaaS layer.

### A market centred on a few ecosystems could affect competition at the PaaS layer

8.64 If the market were to concentrate further towards a few ecosystems with more limited competition between them, then hyperscalers may have less incentives to enable rival ISVs on their platform to attract new customers or maintain existing customers. We have already heard concerns from ISVs that the hyperscalers favour their own first-party PaaS products where they compete with third-party ISV products. We consider it likely that a more concentrated market would increase the ability and incentives of hyperscalers to foreclose or exploit ISVs, for example by acting in ways that favour their own competing products where those compete directly with ISVs.

8.65 This type of conduct could potentially happen through a variety of mechanisms. This could include hyperscalers raising ISV's input costs by increasing the price they pay for access to a hyperscaler's infrastructure, or hyperscalers making interoperability more difficult for ISVs to achieve. If marketplaces were to develop as an essential gateway for ISVs to reach hyperscaler customers in the future, self-preferencing of first-party PaaS products could be another mechanism by which hyperscalers could distort competition at the PaaS layer between their first-party products and those of rival ISVs.

8.66 Rival ISVs may find it harder to monetise their investments, which could result in lower quality or less availability of rival ISV products on hyperscaler ecosystems. This could in turn weaken constraints on hyperscalers' first-party PaaS products leading to higher prices, lower quality and less innovation over time.

# 9.Microsoft's licensing practices

9.1 A number of stakeholders in the cloud computing industry have made submissions to Ofcom regarding the software licensing practices of some cloud providers, in particular Microsoft. Specifically, the submissions allege that the cloud providers in question are using their strong position in software products to distort competition in cloud infrastructure. This, the submissions state, makes it unattractive for customers to use some Microsoft software on non-Microsoft cloud infrastructure. Microsoft disputes the veracity of the practices as alleged in the submissions.[1063]

9.2 The concerns are set out in a number of materials and by a number of parties including the trade association CISPE,[1064] (and reports by Professor Frédéric Jenny which they commissioned (the Jenny reports)),[1065] and the cloud providers [✂][1066] and [✂].[1067] In 2022, CISPE filed a formal complaint against Microsoft with the European Commission relating to these practices.[1068] It has been reported that Microsoft has made an offer to CISPE to settle its concerns, and that these talks are ongoing.[1069] In March 2022, OVHcloud, Aruba, and the Danish Cloud Community filed a complaint with the European Commission against Microsoft. Press reports indicate that Microsoft reached agreement with OVHcloud, Aruba and the Danish Cloud Community resulting the three entities to withdraw their complaints but no such agreement has been reported in relation to the CISPE complaint.[1070] In June 2023, Google in its response to the US Federal Trade Commission's Inquiry into Cloud Computing Business Practices voiced concerns regarding allegedly "unfair" licensing restrictions from some legacy on-premises providers, including from Microsoft.[1071]

9.3 The purpose of this section is to summarise what we have heard and set out briefly in principle what the implications might be for cloud infrastructure competition. We have not undertaken an assessment of the conduct, nor have we made any findings in relation to the concerns themselves.

---

[1063] Microsoft response to the s.174 notice dated 21 October 2022, question 33g, [✂]; and Microsoft response to our follow-up questions dated 27 January 2023 in relation to Microsoft's response to the s.174 request of 21 October 2022.
[1064] CISPE website, 9 November 2022. Executive Summary of CISPE Complaint against Microsoft [accessed 21 September 2023]; and CISPE paper submitted to the EC, 2 May 2023.
[1065] Jenny, Frédéric, October 2021. Cloud Infrastructure Services: An analysis of potentially anti-competitive practices [accessed 21 September 2023]; and Jenny, Frédéric, 21 June 2023. Unfair Software Licensing Practices: A quantification of the cost for cloud customers [accessed 21 September 2023].
[1066] [✂] response to the CFI, [✂]; and [✂].
[1067] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[1068] CISPE complaint to the European Commission, 9 November 2022.
[1069] CISPE website, 20 April 2023. Microsoft opens settlement negotiations with CISPE [accessed 21 September 2023].
[1070] Reuters, 28 March 2023. Microsoft offers to change cloud practices to ward off EU antitrust probe [accessed 21 September 2023]; Bloomberg, 28 March 2023. Microsoft, OVH Prepare to Settle EU Antitrust Complaint on Cloud [accessed 21 September 2023]; MLex, 26 May 2023. Microsoft signs deal with OVH and others to withdraw EU complaints but CISPE probe continues [accessed 21 September 2023]; and The Register, 30 May 2023. Top cloud players reject Microsoft's attempt to settle EU licensing complaint [accessed 21 September 2023].
[1071] Google Cloud's response to the FTC's call for public comment on the FTC's Inquiry into Cloud Computing Business Practices, 21 June 2023. https://www.regulations.gov/comment/FTC-2023-0028-0069 [accessed 21 September 2023].

9.4 The section below focusses on Microsoft. There are also submissions which raise similar issues related to Oracle.[1072] We focus on Microsoft because our study is concerned with cloud infrastructure and Microsoft has a far larger market share compared to Oracle and has grown substantially more than them in recent years, meaning that in principle its conduct has the potential to have a greater effect on competition in the cloud infrastructure market.

# Summary of submissions received

9.5 We have received submissions that Microsoft engages in several practices that make it less attractive for customers to use Microsoft's licensed software products on the cloud infrastructure of a cloud provider other than Microsoft Azure. This could be because it is more expensive[1073] for customers when using Microsoft's licensed products on third-party clouds, or other disadvantages[1074] (such as inability to access some features and reduced availability of security updates compared to running on Azure). The submissions allege that this disincentivises customers from using third-party clouds and impacts on competing cloud providers' ability to compete for customers.[1075]

9.6 The submissions focus on a set of business software products provided by Microsoft, including the Windows Server operating system, Microsoft 365 productivity software suite (also known as Office), Windows 10/11, and Microsoft SQL Server database management system. The submissions largely centre on how Microsoft has changed the way it licences and sells these business software products.

9.7 We note that Microsoft's software licensing practices are complex. Our summary of the practices reflects Ofcom's current understanding based on submissions and materials we have received. It is not designed to comprehensively describe them but simply to enable other stakeholders to have a basic understanding of the submissions and explain how they may be relevant to competition in cloud infrastructure.

## Alleged practices – cost of running software on rival cloud infrastructure

9.8 When Microsoft's customers start moving their existing Microsoft software from on-premises to the cloud, they must ensure they have the rights to use that software in the cloud rather than on premises. There is then a question as to whether their existing on-

[1072] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], pages [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], page [✂]. See also: Jenny, Frédéric, October 2021, page 39, paragraph 81. Cloud Infrastructure Services: An analysis of potentially anti-competitive practices [accessed 21 September 2023].

[1073] See, for example: CISPE complaint to the European Commission, 9 November 2022, page 32, paragraph 127; [✂] response to the CFI, [✂], page [✂]; [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂]; and [✂].

[1074] See, for example: CISPE complaint to the European Commission, 9 November 2022, page 59-60, paragraph 270; CISPE complaint to the European Commission, 9 November 2022, page 40, paragraph 182; [✂] response dated [✂] to our follow-up questions dated [✂] concerning the s.174 notice dated [✂], question [✂], paragraph [✂]; [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 request of [✂], question [✂], page [✂]; [✂]; and [✂] response dated [✂] to our follow-up questions dated [✂] concerning the Ofcom / [✂] meeting, [✂], page [✂].

[1075] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂].

premises software licences give them such rights, or whether they may need to repurchase their licenses or pay for additional usage rights.

9.9 We understand from submissions that Microsoft previously did not charge customers for deploying their software on cloud infrastructure, regardless of the cloud provider. However, in October 2019 Microsoft implemented licensing changes which restricted the ability of certain customers to deploy licences for some software products on non-Microsoft cloud infrastructure.[1076]

9.10 These changes require customers with on-premises software licences to pay additional fees to repurchase their licences for some Microsoft products[1077] if they want to use them on certain cloud providers' infrastructure, despite already owning a licence for the use of the software on premises.[1078] We have heard that this amounts to paying for the same service twice.[1079] We have received submissions that, in some cases, Azure customers do not have to pay these additional fees.[1080]

9.11 One option available to customers wishing to repurchase their licences for some Microsoft software products for use on non-Microsoft cloud infrastructure is Microsoft's Services Provider License Agreement (SPLA) program.[1081] Through the SPLA program, cloud providers can purchase the right to sell some Microsoft software products to customers for use with non-Microsoft cloud infrastructure.[1082] Microsoft sets the wholesale price for products purchased under SPLA and the cloud providers in turn set the retail price to their end customers that incorporate Microsoft products as part of their offering.[1083] Products such a Windows Server, SQL Server, and the Microsoft Office Desktop Applications are available through the SPLA.[1084] A competing cloud provider ([✂]) explained to us that in many instances, the SPLA is – in theory at least – the only alternative option for customers to be able to run a number of Microsoft products on their cloud infrastructure.[1085]

9.12 We have heard from competing cloud providers that purchasing some Microsoft products via the SPLA can be more expensive compared to deploying the same Microsoft software on

---

[1076] [✂] response to the CFI, [✂], page [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], paragraph [✂]; and [✂].
[1077] Some of the products that customers can be required to repurchase licences for include the Microsoft Office productivity apps and Windows Server. We also note that for Windows Desktop, customers can be required to pay additional fees to licence this product again. See: [✂] response to the CFI, [✂], page [✂]; and [✂].
[1078] [✂] response to the CFI, [✂], page [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], page [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], paragraph [✂]; and CISPE complaint to the European Commission, 9 November 2022, page 40, paragraph 180.
[1079] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], page [✂]; [✂] response to the CFI, [✂], page [✂]; [✂]; and CISPE response dated 25 July 2023 to our follow-up questions by email dated 26 June 2023 concerning CISPE's complaint to the EC dated 9 November 2022, page 2, paragraph 10.
[1080] [✂] response to the CFI, [✂], page [✂]; [✂]; and CISPE paper submitted to the EC, 2 May 2023, page 40, paragraph 141.
[1081] [✂] response to the CFI, [✂], page [✂]; and [✂].
[1082] [✂] response to the CFI, [✂], page [✂].
[1083] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂].
[1084] Microsoft document, October 2021. Services Provider Use Rights. Available for download at: Licensing Documents (microsoft.com) [accessed 21 September 2023].
[1085] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂]; and [✂].

Azure.[1086] We have received examples to show this.[1087] One example submitted to us described a potential customer whose Windows server licences would cost over five times the price on a competing cloud infrastructure ([✂]) compared to the price that Microsoft was able to offer for using Azure instead.[1088] In this example, the competing cloud provider ([✂]) told us that it could not match the price Microsoft was able to offer, despite discounting other infrastructure costs significantly.[1089]

9.13 We have been provided with examples to show that the price customers pay to re-license Microsoft software can be a significant proportion of their total contract spend.[1090] In an example submitted to us, if a potential customer had chosen to move their legacy on-premise Windows Server and SQL Server footprint[1091] to a competing cloud provider's cloud infrastructure ([✂]) then over 70% of their total contract costs per year would have been for extra licensing fees arising from Microsoft's software licensing practices.[1092]

9.14 We also received submissions that Microsoft raises SPLA prices on a regular basis, and that these price rises do not apply to Microsoft products hosted on Azure.[1093] We have heard that in some cases SPLA prices have increased by 10-25% per annum,[1094] and for some providers their wholesale Windows Server SPLA prices have increased by over 50%.[1095] One competing cloud provider ([✂]) told us that it absorbed SPLA price increases but eventually had to pass on a small proportion of this price increase to their customers as absorbing this cost was unsustainable in the long-run.[1096] However, another competing cloud provider

---

[1086] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], pages [✂]; [✂] response to the CFI, [✂], page [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], pages [✂]; and [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂].
[1087] [✂]; [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], [✂], page [✂]; and [✂].
[1088] In this example, the customer was looking to move their Windows Server-based on-premises infrastructure to the cloud. The example explains that the reason for this cost differential was Microsoft's licensing policy, which restricted this potential customer from using their existing Windows Server licenses on the competing cloud provider's cloud infrastructure (✂). See: [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], [✂], page [✂].
[1089] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], [✂], page [✂].
[1090] [✂].
[1091] [✂] presentation to Ofcom, [✂], slide [✂].
[1092] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂].
[1093] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], pages [✂]; CISPE complaint to the European Commission, 9 November 2022, page 36, paragraph 157; and CISPE paper submitted to the EC, 2 May 2023, page 48, paragraph 191.
[1094] Jenny, Frédéric, October 2021, page 37, paragraph 71. Cloud Infrastructure Services: An analysis of potentially anti-competitive practices [accessed 21 September 2023]; CISPE complaint to the European Commission, 9 November 2022, page 36, paragraph 157; and CISPE paper submitted to the EC, 2 May 2023, page 40, paragraph 148.
[1095] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂]; [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂], paragraph [✂]; and [✂].
[1096] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂].

([✂]) submitted that it has historically absorbed these cost increases and have not raised prices for customers after Microsoft raised their licensing fees under the SPLA.[1097]

9.15 We received a submission that as a result of Microsoft's licensing practices, only a small proportion of a competing cloud provider's ([✂]) customers run Microsoft products on their cloud infrastructure: they estimate that only <0.5% of their compute service usage runs on Windows 10/11, and an estimated <5% of their compute service usage runs on Windows Server.[1098, 1099]

9.16 Our engagement with customers also suggests that Microsoft's licensing policies are a factor that impacts on customer choice as discussed in Section 6, importance of existing relationships in adjacent software markets.

9.17 In May 2022, Microsoft published a blogpost in which it responded to concerns from certain European cloud providers and announced some changes to the way it licenses its software.[1100] It is our understanding from submissions that as part of these changes, Microsoft allowed customers to use their existing on-premises Microsoft software licences on third-party cloud infrastructure of certain providers.[1101] Customers are eligible to do so if they purchase a 'Software Assurance' subscription[1102] with their existing on-premises Microsoft software licences.[1103]

9.18 We have received submissions that Microsoft has specified a group of cloud providers, called "Listed Providers" (Alibaba, Amazon, Microsoft, and Google), who are not eligible for these changes.[1104] We have heard that this means customers are not able to deploy their existing on-premises licences for certain Microsoft software on Listed Provider cloud

---

[1097] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂], paragraph [✂].
[1098] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], pages [✂].
[1099] We subsequently received additional information from two other stakeholders ([✂] and [✂]) providing estimates of the proportion of their UK public cloud customers that use different Microsoft software products on their public cloud. However, these figures were calculated in different ways, making a like for like comparison difficult in the time and with the data available to Ofcom. See: [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[1100] Microsoft website, 18 May 2022. Microsoft responds to European Cloud Provider feedback with new programs and principles - EU Policy Blog [accessed 21 September 2023].
[1101] As part of these changes, customers can also use new Microsoft subscription licences on eligible third-party cloud infrastructure. See: [✂].
[1102] Software Assurance is a subscription offer that can provide additional functionality and licensing rights associated with the software product that it is purchased with. See: Microsoft website. Microsoft Volume Licensing - Microsoft Software Assurance [accessed 21 September 2023].
[1103] CISPE complaint to the European Commission, 9 November 2022, page 40, paragraphs 180-183; [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂]; and CISPE paper submitted to the EC, 2 May 2023, page 56, paragraph 233.
[1104] CISPE complaint to the European Commission, 9 November 2022, page 40, paragraph 183; [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂]; and CISPE paper submitted to the EC, 2 May 2023, page 45, paragraph 178. See: Microsoft document. https://aka.ms/FlexibleVirtualizationBenefitGuide [accessed 21 September 2023].

infrastructure,[1105] and in these cases are required to pay additional fees to licence these products again.[1106]

9.19 CISPE has previously submitted that these changes have failed to address industry concerns.[1107] Listed Providers remain excluded from these changes,[1108] and CISPE previously explained it is not clear who exactly will qualify to avoid the additional licensing fees that can be required of Listed Providers.[1109] We have also received concerns from stakeholders that there is no restriction on Microsoft adding other competing cloud providers to the list of Listed Providers in the future.[1110]

9.20 We have received submissions that although Microsoft is a Listed Provider, Azure customers are not always affected in the same way by Microsoft's licensing policies.[1111] For example, it is alleged that Azure customers at times do not have to pay additional fees that are required of other Listed Provider customers.[1112]

9.21 A competing cloud provider's submission to us alleges that Microsoft has excluded Azure from the same restrictions as other Listed Providers, and markets this exclusion as the "Azure Hybrid Benefit".[1113] Microsoft offers the Azure Hybrid Benefit to customers moving their Windows Server or SQL Server workloads to Azure.[1114] A Microsoft Azure blog post describes how the Azure Hybrid Benefit means Azure customers "don't pay double": "when using cloud services from other providers, organizations are required to pay for both the infrastructure and the licenses. With Azure Hybrid Benefit, you pay only for additional

---

[1105] [✂] response dated [✂] to our follow-up questions dated [✂] concerning the s.174 notice dated [✂], question [✂], paragraph [✂]; [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂]; and CISPE paper submitted to the EC, 2 May 2023, page 38, paragraph 141.

[1106] [✂] response to the CFI, [✂], page [✂]; [✂]; and CISPE response dated 25 July 2023 to our follow-up questions by email dated 26 June 2023 concerning CISPE's complaint to the EC dated 9 November 2022, page 2, paragraph 10.

[1107] CISPE complaint to the European Commission, 9 November 2022, pages 56-58.

[1108] CISPE complaint to the European Commission, 9 November 2022, page 40, paragraph 183; CISPE paper submitted to the EC, 2 May 2023, page 45, paragraph 178; and CISPE response dated 25 July 2023 to our follow-up questions by email dated 26 June 2023 concerning CISPE's complaint to the EC dated 9 November 2022, page 2, paragraph 6.

[1109] CISPE complaint to the European Commission, 9 November 2022, page 40, paragraph 182.

[1110] CISPE complaint to the European Commission, 9 November 2022, page 58, paragraph 263; CISPE paper submitted to the EC, 2 May 2023, page 47, paragraph 178; CISPE response dated 25 July 2023 to our follow-up questions by email dated 26 June 2023 concerning CISPE's complaint to the EC dated 9 November 2022, page 2, paragraph 12; and [✂] response dated [✂] to the s.174 notice dated [✂], question [✂], page [✂], paragraph [✂]. CISPE also explained that there are still "limitations" to using Microsoft software on non-Azure cloud infrastructure compared to using the software on Azure. See: CISPE paper submitted to the EC, 2 May 2023, page 44, paragraph 177.

[1111] [✂]; CISPE paper submitted to the EC, 2 May 2023, page 55, footnote 119; and CISPE response dated 25 July 2023 to our follow-up questions by email dated 26 June 2023 concerning CISPE's complaint to the EC dated 9 November 2022, page 2, paragraph 8.

[1112] [✂] response to the CFI, [✂], page [✂]; CISPE response dated 25 July 2023 to our follow-up questions by email dated 26 June 2023 concerning CISPE's complaint to the EC dated 9 November 2022, page 2, paragraphs 8-10; and [✂].

[1113] [✂] response to the CFI, [✂], page [✂].

[1114] See: Microsoft Azure Website. Azure Hybrid Benefit - Hybrid Cost Calculator | Microsoft Azure [accessed 21 September 2023].

infrastructure. You will need to repurchase your Windows Server license on other providers' clouds."[1115]

9.22 We have also heard that for Windows 10/11, which is not offered via SPLA, customers using Listed Providers' cloud infrastructure are required to pay additional fees, and at times these fees may not apply to Azure customers in the same way.[1116] A competing cloud provider ([✂]) submitted that these additional fees can result in a significantly higher annual per user cost for Listed Provider customers.[1117]

## Alleged practices – non-cost issues for users of rival cloud infrastructure

9.23 We have also heard that there may be other differences between non-Azure customers and Azure customers when purchasing Microsoft products which may disincentivise the use of non-Azure cloud infrastructure.

9.24 It has been submitted that in some cases, non-Azure customers get an "inferior" version of some Microsoft software products compared to Azure customers because non-Azure customers are not able to receive some services that are available to Azure customers.[1118] The submissions we have received include concerns that Azure customers can access additional features and security updates when using some Microsoft software products on the cloud compared to non-Azure customers.

a) For example, we have heard that Microsoft reserves some newer features of SQL Server,[1119] Windows Server[1120] and Windows Desktop[1121] exclusively for Azure customers, and there can be limitations associated with the security updates for some Microsoft products when using them on non-Azure clouds that do not apply to Azure customers.[1122]

---

[1115] See: Microsoft Azure blog, 26 January 2022. Save big by using your on-premises licenses on Azure | Azure Blog and Updates | Microsoft Azure [accessed 21 September 2023].
[1116] The submissions explain that purchasing a Virtual Desktop Access (VDA) user subscription – used by customers to license Windows for use with virtual machines – is an additional cost for Listed Provider customers that in some cases does not apply when the customer is running Windows 10/11 on Azure. See: [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂]; [✂] response dated [✂] to our follow-up questions dated [✂] concerning the s.174 notice dated [✂], question [✂], paragraph [✂]; [✂]; and CISPE paper submitted to the EC, 2 May 2023, page 37, paragraph 142.
[1117] [✂] response dated [✂] to our follow-up questions dated [✂] concerning the s.174 notice dated [✂], question [✂], paragraph [✂].
[1118] [✂] response to the CFI, [✂], page [✂].
[1119] Microsoft has launched the Stretch Database feature with SQL Server that automatically stretches running databases from on-premises to Azure; this is only available to Azure customers. See: [✂] response dated [✂] to our follow-up questions dated [✂] concerning the s.174 notice dated [✂], question [✂], paragraph [✂].
[1120] There is an Azure edition of Windows Server 2022 Datacenter with additional features that are unavailable on the versions of Windows Server available to non-Azure customers. See: [✂] response the CFI, [✂], page [✂].
[1121] CISPE have also explained that only Azure customers can run the more CPU-efficient Windows 11 multi-session. See: CISPE complaint to the European Commission, 9 November 2022, page 60, paragraph 270.
[1122] [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], [✂], page [✂]; CISPE paper submitted to the EC, 2 May 2023, page 65, paragraph 286; and [✂].

b) We have also heard that customers cannot run Microsoft 365 on Listed Providers' cloud infrastructure.[1123]

## Microsoft's response to summary regarding its software licensing practices

9.25 Microsoft disputes the veracity of the submissions.[1124] When given summary details of the submissions in preparation for the interim report and the opportunity to comment, Microsoft has stated that although Ofcom has not shared specifics of the submissions with them, Microsoft has pointed to the fact that its software is by far a minority use case in the cloud as compared to open source solutions. Moreover, Microsoft has shared its view that competition is robust, Azure is not the market share leader, and many cloud providers are growing at double digit rates.[1125]

# Relevance for competition in cloud infrastructure

9.26 The nature of the submissions that we have received is that Microsoft engages in a number of behaviours in relation to its business software products which encourage customers to use this software on Azure rather than on competitors' public cloud infrastructure. The submissions argue that this conduct leads to an adverse impact on competition in cloud infrastructure services.

9.27 In principle, Microsoft's behaviour in relation to its business software which leads to higher costs and / or technical disadvantages to using that software on rival infrastructure, could have an impact on customer choices. If this were to lead to customers favouring Azure, it could make it more challenging for rival cloud providers to gain customers. This could risk dampening competition if it were to further undermine customers' ability to threaten to switch some or all of their workloads to rivals, and compound the concerns we have provisionally identified in Section 8.

9.28 We have not undertaken an assessment of the submissions or what the exact nature of any impact on competition in cloud infrastructure might be. This would require, among other things, a more detailed understanding of Microsoft's approach to licensing, the importance of Microsoft's software to cloud infrastructure customers (including the presence of any alternative options) and whether rival cloud providers can absorb and respond to any cost and / or technical differences that might exist. Therefore, we are making no findings in relation to the complaints themselves. It will be for the CMA to decide whether to investigate these issues further during the market investigation.

[1123] [✂] response dated [✂] to our follow-up questions dated [✂] concerning the s.174 notice dated [✂], question [✂], paragraph [✂]; [✂] response dated [✂] to our follow-up questions dated [✂] in relation to [✂] response to the s.174 notice dated [✂], question [✂], page [✂]; CISPE paper submitted to the EC, 2 May 2023, page 33, paragraph 121; and [✂].
[1124] Microsoft response to the s.174 notice dated 21 October 2022, question 33g, [✂]; and Microsoft response to our follow-up questions dated 27 January 2023 in relation to Microsoft's response to the s.174 notice dated 21 October 2022.
[1125] [✂].

# 10. Overview of potential interventions

10.1 In our report, we have set out our findings on competition in cloud infrastructure services. We identified several reasons why the market for cloud infrastructure services may not be working well for customers and ultimately UK consumers. While there are inherent technical barriers to switching and adopting more integrated multi-cloud architectures, we identify several features that may further raise these barriers in ways that risk limiting competition.

10.2 This section sets out a high-level overview of the potential merits, risks and challenges of credible intervention options that may address the competition concerns that we have identified in our report. Some may only be fully effective if implemented in conjunction with others or, in some cases, implementing one may mitigate the need to implement another. As with any regulatory intervention, there are a number of potential risks and costs from interventions in cloud infrastructure services. These are likely to differ depending on the exact design of the intervention.

10.3 We remain of the view that there are remedies available which could address the competition issues we have identified in this market study. During the course of a market study, we are not required to assess the most appropriate intervention, rather to identify that there are feasible remedies available which could address our competition concerns. Therefore, we consider some comments on remedies raised by stakeholders below but have not sought to determine whether any individual intervention would be justified or exactly how it might be designed.

10.4 We are not making recommendations or advocating for any specific interventions at this time – this will be for the CMA if it finds an adverse effect on competition in relation to public cloud infrastructure services in the UK and considers that remedies would be appropriate.[1126]

10.5 We are also aware of related regulatory developments, both in the UK with the Digital Markets, Competition and Consumers Bill[1127] and internationally, that have the potential to impact the competitive landscape for cloud infrastructure services. In particular, the EU Data Act aims to introduce requirements relating to egress fees and interoperability which we consider to be comparable to remedy options we explore in this section.[1128] When considering potential interventions in the UK, we are cognisant of proposals elsewhere, especially given the global nature of cloud infrastructure services.

## Improving competitive outcomes by reducing barriers to multi-cloud and switching

10.6 As set out in Section 8, it is our view that whilst there are positive signs of competition in cloud infrastructure services at present, there are a number of barriers to effective

[1126] In Section 11 we set out our reasoning for making a market investigation reference to the CMA.
[1127] See Section2.
[1128] 2022/0047 (COD) Data Act, Legislative Observatory European Parliament [accessed 15 August 2023].

competition. We see indicators of harm today and risks for this to become worse in future, with barriers to multi-cloud and switching at the core of our concerns.

10.7 Reducing barriers to multi-cloud (particularly integrated multi-cloud) and switching (especially switching between clouds) would improve customer outcomes now and reduce the risk of competition deteriorating in the future. In the short-term, it would strengthen the ability of existing customers to move current and future workloads to rival clouds, which in turn would strengthen the competitive pressure placed on cloud providers. This should result in lower prices compared to what would be the case if customers face high barriers to switching and multi-cloud. This should also strengthen the ability for customers to build solutions across different providers that best satisfy their needs.

10.8 Our findings suggest that reducing barriers to multi-cloud and switching could help to limit the risk that the market trends towards greater concentration and further dampening of competition. It could enable greater scope for smaller providers to gain scale by challenging the market leaders for all or some of the workloads of their customers. As set out in our report, we think there are real risks that smaller cloud providers will find it increasingly difficult to expand as the growth of new customers slows, and an increasing number of existing customers face material barriers to switch all or substantial parts of their demand away from the ecosystems of the market leaders – AWS and Microsoft.

10.9 We are concerned that this could weaken the competitive constraints on the market leaders, with barriers to multi-cloud and switching allowing them to entrench their positions and avoid competing vigorously with each other. We also are concerned that this could give them a greater ability and incentive to foreclose rival ISVs – lowering choice, quality and raising eventual prices for customers.

10.10 We think that the combination of barriers is keeping demand for multi-cloud low, which is limiting the development of market-led solutions that can facilitate interoperability between clouds. If barriers to multi-cloud are lowered, demand is likely to respond accordingly. This would provide incentives to encourage innovation and the growth of solutions which can facilitate interoperability between existing clouds, as well as the take-up of cloud-neutral and interoperable technologies.

10.11 At this stage, we have not identified the need for specific interventions solely limited to the use of public cloud in broadcasting or telecoms. As set out in Section 3, we think that the issues these groups face are consistent with the broader market. We would therefore expect any interventions specifically targeted at reducing barriers to multi-cloud and switching to benefit customers in these sectors. However, we recognise that the use of public cloud is still evolving in these sectors, particularly telecoms. We will continue to monitor developments using the evidence and expertise built through this study.

# Overview of potential interventions

10.12 We have identified potential interventions that we think could reduce barriers to multi-cloud and switching. These are set out below, grouped under each of the barriers we identified in earlier sections.

## Data egress fees

10.13 As set out in Section 5, customers often need to transfer their data in order to run their business – whether that is for storage, resilience back-up, or for data processing. Some cloud

providers charge customers for transferring that data out of their cloud infrastructure – we refer to these charges as egress fees. We are concerned that the charging of egress fees has an impact on customer choices, incentivising the purchase of multiple cloud services through the same provider and limiting their ability to operate an integrated multi-cloud architecture, as well as increasing switching costs.

10.14 We discuss below a range of possible remedy options that could reduce the impact of egress fees on customers' ability to multi-cloud and switch.

### Equalise egress fees with other charges

10.15 One approach to mitigating the adverse effects of egress fees could be to require providers to set them in line with other standard data charges incurred by customers. For example, egress fees could be set at a rate which is no higher than the price of internal data transfers within a cloud. This means that customers would face the same cost for data transfer whether they are moving data between clouds or within the cloud of a single provider. The aim would be to help facilitate the take-up of multi-cloud architectures and reduce frictions to switching. This would foster competition based on quality of services and the price of using those services, rather than the price of transferring data to them.

10.16 Stakeholders agreed with our view that determining exactly which other data transfer charges to equalise egress fees against would be complex. Charges vary based on factors such as the volume of data transferred, location of originating and destination data centres and type of infrastructure used to transfer the data (e.g. private network or public internet).[1129] There also may be differences between the cost of providing internal data transfer and external data transfer.

10.17 Once an appropriate data transfer charge is identified, there would be a risk that cloud providers raise this price in line with egress charges, rather than lowering egress fees accordingly.

### Place price controls on egress fees 'at cost'

10.18 A price control that restricts egress fees to at cost charges is likely to reduce the price of egress from current levels, making integrated multi-cloud deployments and switching a more feasible choice. Under this approach, cloud providers would be able to recover the costs of providing the infrastructure and management associated with data transfer from those customers who are moving data. In principle, prices set at cost would ensure that an efficient amount of data transfer occurs. Settling on the right price control level would require careful consideration, and the costs of administering a price control would also need to be factored into an assessment of this intervention.

### Prevent providers from charging for data egress

10.19 The most straightforward way of designing an intervention on egress fees is to prevent cloud providers from charging them at all. Given that not all cloud providers charge egress fees, and an alliance of smaller cloud providers has been founded with the intention of waiving egress fees, this would place the hyperscalers in line with other providers rather than setting a new industry-wide practice.[1130] We note that measures to gradually withdraw switching

[1129] Microsoft response to interim report, page 56-7.
[1130] Cloudflare website. Bandwidth Alliance [accessed 2 August 2023].

charges (including data egress fees) will be introduced in the EU via the forthcoming Data Act.[1131]

10.20 The intention of the remedy would be to make data egress more feasible. This would increase the ability of customers to switch and adopt more integrated multi-cloud architectures, increasing competition between cloud providers. When considering this remedy, these benefits would need to be traded off with the potential for unintended negative effects.

10.21 Several stakeholders raised concerns about unintended consequences, such as resulting in "excess (suboptimal usage)".[1132] Another risk is that service prices might rise where egress fees are capped,[1133] and that this could leave customers who do not transfer much data (which may be small customers) subsidising those which transfer lots of data (which are more likely to be large customers).[1134] We recognise the risk of potential adverse effects, but we do not think a significant waterbed effect on other service prices is likely given that our analysis in Section 5 indicates that:

a) hyperscalers are likely to be setting egress fees above the cost of transfer;
b) hyperscaler revenues from egress are currently a relatively small proportion of their total cloud revenue; and;
c) some large customers already receive large discounts on data transfers.

10.22 Some stakeholders have argued that our definition of egress is overly expansive, because it includes forms of data transfer which are not related to moving data between cloud providers, which could lead to overly expansive, disproportionate remedies which have unintended consequences.[1135] For example, setting prices below cost could lead to more data egress to end users (for example, more streaming of data directly from the cloud to end users) than is socially optimal. This in turn could require providers to build additional infrastructure in response, which could further increase the overall level of fixed costs in the industry.[1136]

10.23 We do not consider that our definition is overly expansive. Given our findings that egress fees act as a barrier to adoption of integrated multi-cloud solutions and that the potential remedies (such as banning egress fees) would address this, and are feasible, we consider investigating this option further does not appear to be disproportionate. In this context, we also note that [✂] have told us they cannot identify the purpose of a customer's data transfer (i.e. whether it is a normal part of that customer's business or a customer transferring data to a rival cloud provider).[1137] We are only at market study stage and we think removing egress fees (either limited in scope, or across the board) remains a feasible

[1131] European Commission. Data Act: Commission proposes measures for a fair and innovative data economy [accessed 11 September 2023].
[1132] Microsoft response to the interim report, page 58.
[1133] AWS response to the interim report, page 9.
[1134] AWS response to the interim report response, page 9.
[1135] Google response to the interim report, page 2; University of East Anglia response to the interim report, page 3.
[1136] AWS response to the interim report, pages 7-8.
[1137] [✂].

remedy option for the CMA to consider if it finds an adverse effect on competition in relation to public cloud infrastructure services in the UK. [1138]

10.24 Hyperscalers have also stated that removing egress fees would undermine providers' incentives to invest in their networks (including innovation incentives).[1139] However, we consider that network investment is unlikely to be driven exclusively by the potential external data transfer fees which can be earned (given the small proportion of revenue accounted for by egress fees) and more likely to be determined by demand for the different services which flow across that network and benefit from network investments. Greater competitive pressure on additional workloads from lower egress fees may increase the incentive for providers to invest in the quality of their network, as the quality of their network is likely to be a factor in being able to attract and retain customers for the range of cloud services they offer.

## Technical interoperability and portability

10.25 As set out in Section 5, we are of the view that technical barriers hinder customers' ability to multi-cloud and switch. Hyperscalers said that they provide services to facilitate multi-cloud and switching, but our analysis suggests these are not sufficient to mitigate the barriers faced by customers. While we recognise that some technical barriers may be inherent, the evidence we have collected during the study suggests that some technical barriers may not be justified, and market forces do not appear to be strong enough to lower them. However, we think there are viable remedies that would have the potential to mitigate these barriers.

10.26 In response to our CFI and interim report, we received several recommendations on different measures related to technical interoperability and portability of cloud infrastructure services, which we have taken into account when considering the feasibility of potential remedies.[1140] We have broadly grouped the potential interventions into four categories:

a) Requirements for suppliers of cloud services to be more transparent about the interoperability of their cloud services.

b) Requirements for suppliers of cloud services to make their cloud services easier to interoperate with.

c) Requirements for suppliers of cloud services which aim to increase the degree of standardisation.

d) Requirements for cloud providers to interconnect their data centres.

### Requirements for suppliers of cloud services to be more transparent about the interoperability of their cloud services

10.27 We have seen indications that hyperscalers may not be transparent about the degree of interoperability of their cloud infrastructure services or that their published documentation is not always sufficiently clear. Inadequate transparency and documentation can serve as a

[1138] We note that the EU's Data Act intends to withdraw egress charges related to switching (following a transition period of three years). In principle therefore, it does not appear to be a barrier to concluding that the remedies considered would not be feasible.
[1139] AWS response to the interim report, paragraph 18, page 9; and Google response to the interim report, paragraph 9, page 2.
[1140] Vodafone response to the CFI, page 4; Cloudflare response to the CFI, page 1-3; and [✂].

barrier to multi-cloud and switching, as well as integrating first-and third-party services on the same cloud. For example, customers may build a cloud architecture which they believe to be highly interoperable, without realising that they are relying on different cloud technologies.

10.28 A potential intervention to address such concerns would be to require the hyperscalers to publish documentation on the interoperability of their cloud infrastructure services. This would include, for example, requiring hyperscalers to explain the compatibility of their cloud infrastructure services with open-source software. We believe this would allow customers to make more informed choices when designing their cloud architectures and facilitate the integration of multiple clouds and switching between clouds. It would also facilitate the integration of first- and third-party services (i.e. multi-vendor cloud architectures) and switching within the same cloud.

10.29 While hyperscalers commented on their current levels of transparency (discussed in Section 5 and 7), we did not receive substantial stakeholder submissions on the appropriateness of a potential transparency remedy. We have nevertheless considered the potential trade-offs with such an intervention. On the one hand, we would expect there to be low implementation costs as it would simply entail the publication of additional information rather than adjusting any existing systems. On the other hand, an extensive transparency requirement in a dynamic market might create an unnecessary burden on hyperscalers. For example, if they have to immediately reflect in their documentation the implications of any change in the design of their cloud infrastructure services.

10.30 In the interim report we said that another route available to us as part of the market study could be to publish advice to customers on how to build their cloud architectures in a way that keeps them as flexible as possible. We have decided not to do so at this point in time. This is because there is already information available[1141] and so the incremental impact of publishing additional advice now would likely be limited. We also consider that such advice may not be sufficient to change customer behaviour in a way that truly fosters competition. For example, advising customers to be more cloud-agnostic by deploying additional abstraction layers and increasing their use of adaptors can significantly increase customers' costs (such as re-engineering fees, egress fees, and subscription costs for additional services) which may result in limited changes in actual behaviour.

### Requirements for suppliers of cloud services to make their cloud services easier to interoperate with

10.31 As set out in Section 5, our analysis indicates that AWS and Microsoft limit the interoperability of some of their cloud services with third parties. In particular, AWS and Microsoft design some of their cloud infrastructure services such that certain functionalities are only accessible by other first-party cloud services (i.e. asymmetry of functionalities). As set out in Section 7, even where access is not restricted, there may be a significant delay in these functionalities becoming available to third-parties due to the time they need to update their cloud services after the cloud provider publishes an updated API. This limits customers' ability to integrate multiple clouds where AWS and Microsoft cloud infrastructure services only interoperate with their respective clouds. It also limits customers' ability to switch between clouds because they are not able to separate out the service they want to keep

[1141] For example, see IT Strategy Template for a Successful Strategic Plan | Gartner [accessed 27 September 2023], Provisioning a Multi-tenant CSP Agnostic Cloud Platform for the Federal Government (hashicorp.com) [accessed 27 September 2023], and What is cloud agnostic? | VMware Tanzu [accessed 27 September 2023].

(e.g. a PaaS service) from that which they might want to switch (e.g. an IaaS service). To a lesser extent, this also limits customers' ability to integrate and/or switch between first- and third-party cloud services hosted on the same cloud.

10.32 In addition, as set out in Section 7, we have also heard that the hyperscalers may be unfairly favouring their own services. Some stakeholders have provided examples of cases where AWS and Microsoft allow limited access to their proprietary APIs, thereby reducing the discoverability and ease of use of third-party cloud services. This would potentially affect the ability for rival ISVs to compete with the hyperscalers' first-party cloud services. This includes ISVs that offer cloud-agnostic services that are important for customers' ability to multi-cloud and switch.

10.33 A way to address these concerns would be to set requirements that allow third parties to interoperate with individual elements of AWS and Microsoft's services to the same extent as AWS and Microsoft currently do when providing their own cloud services. This approach would set outcomes rather than define any aspects of the technical design of hyperscalers' services. We have identified two types of requirements that could be put in place.

*A service access requirement*

10.34 First, a service access requirement for AWS and Microsoft to make the individual elements of their cloud infrastructure services available to use in conjunction with those of third parties. For example, this would allow a customer to purchase an AWS PaaS service and use all of its functionalities in conjunction with the IaaS service of a rival cloud provider. This would increase customer choice and require AWS and Microsoft to facilitate interoperability with their cloud infrastructure services, thereby lowering the costs of integrating multiple public clouds and switching between clouds. It would also facilitate the integration of first- and third-party services (i.e. multi-vendor cloud architectures) and switching within the same cloud.

10.35 However, the requirement to make services available might only guarantee a basic level of technical interoperability. Customers may still find integration is more difficult, quality of experience is reduced, or functionalities are becoming available more slowly when using first-party and third-party services in combination, by comparison to using an integrated service provided solely by AWS or Microsoft.

*A complementary equivalence requirement*

10.36 This means a second requirement around equivalence may be needed to ensure AWS and Microsoft's cloud infrastructure services (for example a type of PaaS service) can interoperate with third-parties (for example IaaS service from another cloud provider) in the same way as first-party cloud services. In practice, this could be achieved by a requirement for AWS and Microsoft to provide equivalent access to their cloud infrastructure services (for example, by designing the APIs associated with them to support third-party inputs). This would also limit the ability of AWS and Microsoft to circumvent the access requirement by making functionalities available to third-parties, but with only a limited degree of interoperability or with a delay. This could reduce technical barriers to multi-cloud, potentially increasing the demand for multi-cloud and switching, which could incentivise market-led solutions that can facilitate interoperability between clouds. It could also facilitate the integration of first- and third-party services (i.e. multi-vendor cloud architectures) and switching within the same cloud.

10.37 In its response to the interim report, a cloud provider [✂] said that mandating exposure of internal functions could (i) have serious detrimental effects on the integrity, security and quality of services provided, (ii) increase the costs to the customer, and (iii) significantly decrease the pace of innovation. On the first point, the cloud provider said that, since changes to customer interfaces must happen more slowly to avoid disruption of customer workloads, treating every internal interface as "public" would seriously impede the rate at which it can improve its services' functionality and security. Further, the cloud provider also said that an intervention would also create an operational burden – from the volume of traffic alone – that would dramatically increase the difficulty of effectively monitoring and updating all public APIs, increasing the potential for bad actors to find vulnerabilities and carry out attacks. On the third point, the cloud provider said that exposing all of a service's internal interfaces could allow competitors to reverse-engineer various functions and features, and discover their underlying structure, data flow or logic. It said this proposal would likely erode incentives to invest in research and development as any successful innovation would be exposed to competitors. It also said that, aside from reducing the range of services and features available to customers, mandating fully open internal functionality would also adversely impact smaller or newer cloud providers whose key selling point to customers might be the offering of unique functionalities. Finally, this stakeholder said that cloud providers need the ability to make the practical decisions on what internal functionality to make public based on the customer demand, technical and security implications, and the resources available.[1142]

10.38 In relation to integrity, security and quality of services provided, we acknowledge that designing these interventions would be complex. However, we believe that these challenges could be overcome through carefully defining the exact scope of the interventions, including whether they would apply to specific functionalities, services or suppliers.

10.39 In relation to costs and innovation, we acknowledge that these interventions could lead to implementation costs and potentially dampen incentives to innovate. However, cloud providers would still be able to monetise individual services and functionalities, which should preserve their incentives to innovate at service-level.[1143] Moreover, the potential for smaller providers to 'wait in the wings' with innovative solutions that can work in conjunction with the services of existing cloud providers is important to ensure lasting constraints on market leaders.

10.40 The nature and impact of these interventions would require careful consideration. However, we do not believe that the challenges associated with them are insurmountable. In particular, we note that other suppliers of cloud services are already developing their cloud services to be more interoperable (e.g. Google, IBM, Oracle and a number of ISVs).

## Standardisation of cloud technologies

10.41 In Section 5 we described how the technical differentiation of cloud technologies associated with the interoperability and portability of cloud infrastructure services and ancillary cloud services could be a barrier to multi-cloud and switching. These technologies include APIs, protocols, workflows, programming languages, data formats or other technologies the customer cloud architecture relies on. While there are services available in the market to

[1142] [✂].

[1143] For example, Microsoft said that it generates revenue every time its APIs are called, even from a different cloud architecture (Microsoft response to the interim report dated 31 May 2023, paragraph 209).

assist customers to mitigate the lack of standards through various means, there appears to be limitations to their effectiveness.

10.42 Setting technology standards has led to increased interoperability and portability in various sectors, for example in telecoms. While there have been some industry-led standardisation efforts in cloud, our assessment to date is that there has not been a single concerted approach to standardisation and many attempts have been unsuccessful in gaining industry-wide traction.

10.43 One intervention would be to use voluntary standards to improve interoperability and portability in cloud technologies. A starting point would be to generate support for existing industry standards and open-source software, which may help improve their availability and quality. However, using a voluntary approach means the incentives of hyperscalers, and particularly AWS and Microsoft, are unlikely to align with the use of established standards unless there is significant industry pressure to do so.

10.44 Another approach would be to mandate the use of specific standards, which would guarantee broader adoption. We have assessed how standards could apply in two areas: cloud infrastructure services and ancillary cloud services.

*Mandatory standardisation of cloud infrastructure services*

10.45 In relation to cloud infrastructure services, mandating standards may address some of their technical differentiation, including any asymmetry of functionalities. But it may also come with significant complexity and risks of unintended consequences. The complexity includes deciding which services should be captured by any standards, defining their technical aspects such that any standard is sufficiently flexible and technology neutral, and deciding how they would apply across the different underlying clouds. The unintended consequences include significant implementation costs if standards require re-engineering of cloud services and customer applications, and a risk to innovation if standards become outdated.

10.46 Stakeholders that commented on mandating standardisation as a potential intervention raised concerns that it may lead to a significant risk to innovation. In particular, stakeholders said that mandatory standardisation of cloud infrastructure services may result in suppliers limiting their innovation to the lowest common denominator.[1144] One cloud provider [✂] added that standardisation will make it harder for smaller cloud providers to differentiate themselves and thereby compete. It also said that regulatorily-enforced standardisation is incompatible with dynamic and innovative industries, such as the IT sector.[1145]

10.47 We broadly agree with these stakeholders' submissions. In relation to the potential impact on innovation, we are concerned that mandatory standards may impact the ability of suppliers to offer new solutions, services and functionalities that address customers' existing and future use cases.

*Mandatory standardisation of ancillary cloud services*

10.48 In relation to ancillary cloud services, mandating standards may reduce the difficulty that customers currently experience in managing multiple clouds. In this way, such a potential

[1144]For example, [✂]; [✂] confidential response to the interim report, pages [✂]; [✂] confidential response to the interim report, question [✂]; CCIA response to the interim report, question 8.5; TechUK response to the interim report, page 3; University of East Anglia response to the interim report, page 6; and Microsoft response to the interim report, paragraphs 28(ii) and 222-223;
[1145] [✂].

intervention would facilitate integration of multiple clouds and switching between clouds. It would also reduce ISVs' costs to making their services available on multiple clouds. Unlike cloud infrastructure services, mandating standards on ancillary cloud services may not be as complex, and implementation costs and the risk of unintended consequences may be lower.

10.49 In relation to implementation costs, ancillary cloud services typically represent a small proportion of the product portfolio of a cloud provider. Therefore, we expect any re-engineering to be a relatively low cost for them. However, there is still a need for careful consideration of the impact of such an intervention, particularly on smaller cloud providers.

10.50 Stakeholders that commented on this potential intervention acknowledged that standards may limit innovation in ancillary cloud services but noted that in this area innovation may be more limited or less important compared to the potential benefits for customers and ISVs.[1146]

10.51 We agree with these stakeholders' submissions on innovation. In particular, we consider that ancillary cloud services generally aim to provide deeper integration with customers' cloud architectures, and harmonisation of these services will not inhibit the ability of cloud providers to continue developing innovative products and functionalities that address customers' existing and future use cases in the cloud.

### Connecting data centres

10.52 As discussed in Section 5, one technological barrier to multi-cloud and switching is data gravity. In particular, we said that latency is one key factor that can contribute to creating a data gravity effect such that the cloud where the bulk of a customer's data is hosted is likely to attract more of this customer's data, as well as associated customer applications and cloud services.

10.53 We understand that physically interconnecting data centres of different cloud providers is likely to be a technically feasible solution to help mitigate such cross-cloud latency, particularly when the data centres are located within the same availability zone.[1147] With a direct interconnection between clouds, data traffic remains on a private network which can increase the speed of data movements, enhance control of network traffic and reduce the chance of unexpected increases in latency. This would help customers, especially those that run latency sensitive workloads, to more easily integrate multiple clouds and switch between clouds.

10.54 We expect such intervention to require a complex design, including which cloud providers it applies to, defining an availability zone and setting out requirements on the type of connectivity. The intervention may also lead to a number of unintended consequences, including high implementation costs, especially for smaller cloud providers, and potentially impact the competitive dynamics in the UK business connectivity market.

10.55 We nevertheless consider that such intervention may on balance have positive effects on competition and customer welfare if designed appropriately.

---

[1146] For example, Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]; Ofcom / [✂] meeting, [✂]; and Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂]. Also, see Eclipse Foundation, Unlock cloud interoperability to foster the EU digital market, April 2023, available at Unlock the Cloud Interoperability to Foster the EU Digital Market (eclipse.org) [accessed 11 September 2023].

[1147] Physically interconnecting data centres may also help overcome other challenges such as data security.

## Committed spend discounts

10.56 As explained in Section 5, committed spend discounts involve a customer agreeing to spend a minimum amount on a cloud provider's products over a period of time, receiving a percentage discount in return. Such discounts are widely used in the cloud market, including by all three hyperscalers, primarily being privately negotiated with the largest cloud customers – a high proportion of whom have substantial spending commitments to hyperscalers.

10.57 We have seen indications that the structure of these privately negotiated cross-service committed spend discounts, when used by hyperscalers, may create an additional incentive for large customers to concentrate all or most of their cloud spending with a single hyperscaler, even for workloads where there are few technical barriers to using multiple providers.

10.58 In our interim report, we emphasised the need for any remedies focusing on the structure of hyperscalers' committed spend discounts to be designed carefully.[1148] Price discounts are important to the customers who have negotiated them, in some cases saving them large amounts of money compared to purchasing services at list prices. Furthermore, the hyperscalers have submitted that commitments from customers give them greater certainty of future demand and so protect investment and innovation.

10.59 In response to our interim report, some stakeholders submitted that any remedy which may affect the discounts received by customers must be designed with great care in order not to dampen the positive aspects of discounts.[1149] A customer submitted that they would be very concerned about interventions which lessened their ability to secure discounts based on their bargaining power, and the regulatory focus should be on ensuring providers cannot impose price increases or impose obstacles on renegotiating customers.[1150] Another customer submitted that they understand the concerns around committed spend discounts, but think other interventions will have a greater impact on competition, specifically those focused on interoperability and egress fees.[1151] A customer which responded to our interim report suggested remedies that improved transparency of the negotiation process (including the factors which hyperscalers consider in making offers to customers) could have a positive impact on competition. [1152]

10.60 The protection of customers who may be unable to use their bargaining power to protect themselves from an expectation of commitment growth (i.e. the pressure to increase their commitment over time) is another possible target for intervention. Some stakeholders suggested that renegotiating customers should have the right to maintain their existing commitment and level of discount.[1153] Such an intervention may have the effects of both preventing pressure to grow commitments, and may increase the bargaining power of customers as they would not face the risk of losing their discount entirely if they fail to accept a hyperscaler's offer before the expiration of their existing agreement.

---

[1148] Cloud services market study (interim report) paragraphs 8.43-8.46.
[1149] For example, UKCTA response to the MIR consultation, paragraph 7.
[1150] [✂] confidential response to the interim report, question [✂].
[1151] BT Group response to the interim report, page 3.
[1152] [✂] confidential response to the interim report, page [✂].
[1153] For example, [✂] confidential response to the interim report question [✂].

10.61 Hyperscalers, in their responses to our interim report, submitted that any interventions focused on committed spend discounts would ultimately be harmful.[1154] They pointed to their submitted rationales for the use of committed spend discounts, that they allow cloud providers to better forecast future customer demand and so lead to higher levels of investment[1155] and stated that interventions focused on them would ultimately raise prices.[1156] Hyperscalers also said that smaller cloud providers require the use of a wide range of discount structures and other tools to compete effectively in the cloud market.[1157]

10.62 We noted that interventions focused on committed spend discounts could result in unintended consequences in our interim report.[1158] Any intervention would need to be targeted at addressing the structure of the discounts that risk distorting competition. It would be important to preserve the ability of cloud providers to gain the commitments of customers to the extent that these are necessary to protect investment and innovation, and also the ability of customers to exercise their bargaining power to gain lower prices and other concessions from cloud providers.

## Challenges predicting cloud spend

10.63 Some customers told us about the challenges they faced when trying to predict their cloud spend, including difficulties associated with understanding pricing and billing processes, challenges in predicting future workloads, and some issues when using cloud monitoring tools. In Section 5 we highlighted the breadth of these challenges.

10.64 We are aware of industry-led approaches to addressing these concerns that could have a beneficial effect on competition. For example, [✂] acknowledged that they need to simplify their pricing structures for customers and provide sellers with appropriate material and tools to explain pricing structures.[1159] Similarly, the development of cost control tools by each hyperscaler demonstrates a positive step to help customers tackle the complexity of cloud pricing and billing, although our evidence suggests that these tools are not without fault themselves.

10.65 We acknowledged in our interim report that these are issues which concern customers and in principle it might be possible through regulatory intervention to increase transparency and simplicity of pricing and billing, for example by imposing a standard approach. We think there are fewer unintended consequences when imposing standards in these areas, as it would maintain flexibility for providers to innovate the underlying product offering. However, we recognise that there are likely to be practical challenges with interventions, given the complexity and diversity of cloud services, and that there is a chance that it could impede innovation by providers in how they price their services.

10.66 In response to our interim report stakeholders suggested a number of possible interventions to issues around predicting cloud spend. Some suggestions were more drastic including an "open, transparent marketplace, akin to the energy market" which would offer businesses

[1154] Microsoft response to the interim report, paragraph 28(i). [✂]. [✂] confidential response to the interim report, paragraph [✂].
[1155] These rationales are assessed in Section 5.
[1156] [✂]. Microsoft response to the interim report, paragraph 174.
[1157] [✂]. [✂] confidential response to the interim report, paragraph [✂].
[1158] Ofcom, 2023. Interim report, paragraph 8.46.
[1159] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].

full price transparency on digital resources.[1160] Others offered simpler suggestions based on simplifying and standardising pricing structures or improving cloud cost monitoring tools.[1161] The consensus among respondents was that making it easier to understand pricing and cloud estimation tools would benefit users and should be the goal.

## Skills

10.67 As set out in Section 5, cloud staff needing to possess specific skills to work on different cloud environments can increase switching costs, as a company wanting to switch or add a new cloud provider may need to retrain or hire staff. We are concerned that while many organisations are putting considerable effort into supporting customers to develop their skills, these costs may be a barrier many customers face.

10.68 The evidence from customers, intermediaries and cloud providers indicates that the time and cost required to address a lack of appropriate in-house skills may act as a strong barrier to multi-cloud and switching. One respondent to our interim report noted the importance of skills in the decision to select a hyperscaler and whether or not to multi-cloud, since training talent continuously and/or for multi-cloud requires high investment and creates cultural challenges.[1162] Another respondent emphasised the finding in our market research that skills issues were a barrier to switching.[1163] Respondents to our market research who had considered switching but had not gone through with it cited the need to retrain staff as one of the most important challenges.

10.69 While Microsoft did not appear to recognise a skills gap (citing that companies generally understand what they are purchasing, and the potential risks involved), others noted the range of measures they are taking to help address the problem. AWS cited a range of training and courses offered by a range of cloud providers and noted their own range of offerings to customers.[1164]

10.70 We expect that reducing technical barriers between clouds will indirectly reduce the skills gap. For example, increased interoperability is likely to make it easier to be trained in multiple clouds. Furthermore, there are wider Government initiatives in relation to digital skills that could more directly contribute to addressing the concern in cloud.[1165]

# Conclusions

10.71 We are not making a recommendation on any specific remedial interventions at this stage. However, our view is that there are credible interventions available that could address the different barriers to effective competition that we have identified through this market study.

---

[1160] Federation of Communication Services response to the interim report, page 5 and Sustainable Digital Infrastructure Alliance response to the interim report, page 4.
[1161] Federation of Communication Services, response to the interim report, page 5, and [✂] confidential response to the interim report, page [✂].
[1162] Priyank Chandra, response to interim report, paragraph 4.1.
[1163] University of East Anglia response to interim report, page 5.
[1164] [✂].
[1165] 2022 Digital Strategy [accessed 11 September 2023]; The UK Science and Technology Framework [accessed 11 September 2023]; Independent Review of The Future of Compute: Final report and recommendations [accessed 11 September 2023].

10.72 We recognise, based on our analysis, that some of the interventions described above have the potential to incur costs or lead to unintended consequences. For example, where they require substantive changes to services or could impact providers' incentives to innovate. There are also likely to be links between different interventions, either where they work better acting in combination or where implementation of one intervention removes the need for another. We note stakeholder concerns on these issues and note that these points will be considered in the round during the CMA market investigation.

# 11. Making a market investigation reference

11.1 In this market study, we have been looking at whether any feature of the market could dampen competition between providers of public cloud infrastructure services and therefore have an adverse effect on customers through higher prices, lower quality products or less innovation.

11.2 We have set out our findings in this report. Our assessment is that, while there are some positive signs of competition at present, there are also clear indications that the market is not working well in certain respects. Our view is that there are features of the market that act as barriers to multi-cloud and switching and we have reasonable grounds to suspect that these features prevent, restrict, or distort competition in the UK, and that they merit further detailed assessment by way of a market investigation reference (MIR).

11.3 Having carefully considered the evidence set out in this report together with the submissions received from stakeholders throughout our market study, we have decided to refer the UK public cloud infrastructure services market to the Competition and Markets Authority (CMA). This section sets out our decision which is informed by the entirety of our final report.

## Our competition concerns

11.4 Public cloud infrastructure services underpin the way we live our lives, run our businesses, and access our public services. Without the cloud, many digital businesses providing services to end consumers would not be able to function in the way they do today. Cloud computing supports not only Ofcom's core sectors, but most other sectors of the economy, for example manufacturing, retail, hospitality and financial services, as well as public and voluntary sector bodies.

11.5 'Public cloud' is the most common cloud deployment model, where cloud services are open to all customers willing to pay, and computing resources are shared between them. However, the public cloud infrastructure services market is relatively concentrated in the UK, with Amazon Web Services (AWS) and Microsoft accounting for the majority of UK revenues in 2022, and Google the main challenger. Taking on board the evidence available, including information from cloud providers and submissions from stakeholders, our report explores and identifies a range of competition concerns.

11.6 Following our assessment, as set out in detail in this final report and its annexes, we have reasonable grounds to suspect that the following features of the market for public cloud infrastructure services prevent, restrict, or distort competition in the UK:

a) Egress fees for moving data out of a provider's cloud;
b) Technical barriers, including, but not limited to, restrictions on interoperability and portability; and
c) Committed spend discounts.

# Scope of the proposed market investigation reference

11.7 As set out in the terms of reference in Annex 6, the MIR covers the supply of public cloud infrastructure services in the UK. For the purposes of this reference:

a) 'Cloud infrastructure services' means services that provide access to processing, storage, networking, and other raw computing resources (often referred to as infrastructure as a service, IaaS) as well as services that can be used to develop, test, run and manage applications in the cloud (often referred to as platform as a service, PaaS).
b) 'Public Cloud' means a cloud deployment model where cloud services are open to all customers willing to pay, and computing resources are shared between them.

11.8 The basis of this scope is explained further below within our explanation of our decision to make a MIR.

# Our decision to make a market investigation reference

11.9 Ofcom has concurrent functions pursuant to section 370 of the Communications Act 2003, with the CMA, under Part 4 Enterprise Act 2002 (EA02) (Market Studies and Market Investigations), with some exceptions. Pursuant to those functions, Ofcom may decide to make a MIR to the Chair of the CMA for the constitution of a CMA Group when we have reasonable grounds for suspecting that a feature or combination of features of a market or markets in the UK prevents, restricts, or distorts competition, and a MIR appears to be an appropriate response.[1166]

## The legal framework

11.10 As set out above, the reference test is one of 'reasonable grounds to suspect'. Where the reference test is met, Ofcom can exercise its discretion to make an MIR. The Competition Appeal Tribunal has recently confirmed that Ofcom's discretion to make a reference is wide and, provided Ofcom has addressed matters sufficiently, that the "reasonable grounds for suspecting" threshold is low.[1167]

11.11 The CMA's guidance on making MIRs sets out four criteria which we have used to help to guide the exercise of our discretion:

a) The scale of the suspected problem, in terms of its adverse effect on competition, is such that a reference would be an appropriate response.
b) There is a reasonable chance that appropriate remedies would be available.
c) It would not be more appropriate to address the concerns through undertakings in lieu of a reference (UILs).

[1166] Section 131(2) of the EA02 sets out what is to be construed as a 'feature' for the purposes of Part 4 of EA02.
[1167] See the explanation of the Competition Appeal Tribunal in *Association of Convenience Stores v OFT* [2005] CAT 36, paragraph 7. See also more recently, *Airwave Solutions Limited & Others v CMA* [2022] CAT 4 at [9]-[10], [12] and [27] and *Apple Inc & Others v CMA* [2023] CAT 21 at [39] were the Tribunal referred to the trigger of that threshold as "low" and one that needs to "viewed in the round".

d) It would not be more appropriate to address the competition problems through alternative powers available to the CMA or through the powers of sectoral regulators.[1168]

11.12 In considering these factors, Ofcom recognises that an MIR leads to significant costs, both to the CMA (and the public purse) and to the parties involved.

## The reference test

11.13 In making an MIR, Ofcom must specify the description of goods or services to which the feature or combination of features concerned relates. However, as explained in the statutory guidance and stated in the CMA's published guidance on the making of MIRs, Ofcom is not obliged to provide a precise definition of the market or markets to which any MIR relates. This is informed by s.131(2) EA02 which explains that features of a market for goods or services are to be construed as: (a) its structure (or any aspect of its structure); (b) the conduct of persons supplying or acquiring goods or services who operate within it; and (c) the conduct of such persons' customers.

### The market

11.14 Our market study assessment has focussed on 'public cloud.' We have also concentrated on 'cloud infrastructure services' which means services that provide access to processing, storage, networking, and other raw computing resources (often referred to as infrastructure as a service, IaaS), as well as services that can be used to develop, test, run and manage applications in the cloud (platform as a service, or PaaS). These are the foundational elements of the cloud stack on which other cloud services (like software as a service, SaaS) are built, and where we currently see the greatest concentration of supply and factors that pose a risk to effective competition.

11.15 The main suppliers of cloud services in the UK are AWS, Microsoft and Google, which provide a full range of cloud services at scale. There are a number of smaller suppliers of cloud services, some offering a broad range of cloud services, while others are more specialised.

### The features of the market

11.16 Based on the evidence and the analysis set out in our report we have reasonable grounds for suspecting that the following features, alone or in combination with each other, prevent, restrict, or distort competition in the supply of public cloud infrastructure services in the UK:

a) Egress fees for moving data out of a provider's cloud. The hyperscalers set these significantly higher than most other cloud providers and the cost of egress fees can discourage customers from using services from more than one cloud provider or switching.
b) Technical barriers, such as, but not limited to, restrictions on interoperability and portability. We have found that customers face technical barriers to interoperability and portability which limits their ability to adopt some multi-cloud and multi-vendor architectures and to switch. While some of the technical differentiation which underpins the barriers may be the result of innovation to the benefit of customers, we are concerned that in some cases such justification may be less clear.

[1168] Guidance about the making of references under Part 4 of the Enterprise Act, OFT 511, paragraph 2.1.

c) Committed spend discounts. The structure of discounts can incentivise customers to use a single hyperscaler for all or most of their cloud infrastructure needs. This can make it less attractive to use rival providers as new needs emerge or to move existing workloads to alternative providers.

11.17 Although we have identified these particular features, our MIR does not confine the CMA's market investigation to these features, and it will be for the CMA to decide whether to investigate other issues during the market investigation.

### Our view on the reference test

11.18 For the reasons set out above in this section and in our report, we have reasonable grounds to suspect that one or more features (alone or in combination) in relation to the supply of public cloud infrastructure services prevent, restrict, or distort competition in the UK and that the reference test is met.

11.19 Having reached this view, we now go on to consider the factors relevant to the exercise of Ofcom's discretion to make an MIR.

## Views on the appropriateness of a reference

11.20 Many responses to our interim report and MIR consultation either supported the reference (explicitly or implicitly) or expressed concern about a specific issue.[1169] In some cases this was a general agreement, support for the reference, or support for intervention in relation to the features we had identified.

11.21 Google was against a reference but considered that if unilateral behaviour was limiting competition a standalone Competition Act 1998 investigation would be more appropriate.[1170] Microsoft and AWS felt a reference was not warranted. Microsoft argued that the market was fundamentally competitive as evidenced by innovation and investment levels, price trends and further expected growth in a dynamic market.[1171] AWS argued that competition for the global provision of IT services, including cloud services, was functioning effectively. In AWS's view there is intense competition, innovation, a variety of providers and good market outcomes. It took the view that there were no grounds on which a MIR could be made.[1172]

11.22 They were supported by a minority of other respondents who had concerns about whether appropriate interventions were likely to be identified following an investigation, including [✂], [✂] and academics from the University of East Anglia.[1173]

---

[1169] BT Group response to the interim report and MIR consultation, page 1; Name Withheld 1 response to the MIR consultation, paragraph 3.3; Oracle response to the MIR consultation, page 3; Name Withheld 2 response to the MIR consultation, paragraph 3.3; Virgin Media O2, response to interim report and MIR consultation, page 1; Federation of Communication Services response to the interim report, page 7; Gener8 response to the interim report, paragraph 5.1 and 8.1; Sustainable Digital Infrastructure Alliance e.V response to the interim report, paragraph 8.1, 8.2 and 8.3; Vodafone response to interim report, page 3; Cloudflare response to the interim report, page 1-3; [✂] response to the interim report, page[✂]; [✂] response to the interim report, paragraph [✂]; [✂] response to the interim report, paragraph [✂]; and [✂] response to the interim report, paragraph [✂].

[1170] Google response to the interim report and MIR consultation, page 4.

[1171] Microsoft response to the interim report and MIR consultation, page 4 and 10.

[1172] AWS response to the interim report, page 14; and [✂].

[1173] University of East Anglia response to the interim report, page 2; [✂] response to the interim report, page [✂].

11.23 However, we have demonstrated in our report that, while there are some good market outcomes, there are features of the market that give us reasonable grounds for suspecting that competition is not working effectively in the UK cloud infrastructure services market.

11.24 While our view is that the legal threshold has been met for a market investigation into the public cloud infrastructure services market, we recognise that an MIR can impose a burden on both the businesses concerned and the CMA. Indeed, Google took the view that an MIR would be inappropriate and place a disproportionate burden on smaller cloud services providers.[1174] It is, therefore, important that we assess the potential significance of the adverse effects on competition that we have reasonable grounds to suspect exist, and that we satisfy ourselves that an MIR is the most effective regulatory response to those effects by considering the scale of the problem and the availability and appropriateness of remedies.

## First criterion: scale of the suspected problem

11.25 In determining the scale of the suspected problem, the CMA guidance identifies three factors of particular significance:

a) The size of the market;
b) The proportion of the market affected by the features; and
c) The persistence of those features.

*Size of the market*

11.26 Cloud services are increasingly important to many businesses and organisations across the economy. We estimate that between 2019 and 2022, UK revenues for IaaS and PaaS combined grew by 35% to 40% per year, and we expect it to continue to grow in line with the global market for cloud services.[1175] In 2022 we estimate UK revenues to be £7.0 billion to £7.5 billion.

11.27 Most UK businesses are at some stage of modernising their IT through the adoption of cloud services. Our customer research[1176] showed that 43% of current users of IaaS/PaaS services expected to be migrating more workloads to the cloud over the next 18 months. When asked about their expectation of change in cloud spend in the next 18 months, 79% of respondents said they expected it to increase slightly or greatly. IDC estimates that around 30% of UK businesses' IT spending went to cloud services in 2022, and it projects this share to rise to 51% by 2027.[1177]

11.28 The health of the cloud market affects the health of an increasing number of sectors. Cloud computing supports not only the communications sector, but other sectors, for example manufacturing, retail, hospitality and financial services, plus public and voluntary sector bodies. Cloud infrastructure acts as a fundamental building block for a diverse range of software applications that ultimately benefit consumers and businesses across the economy.

*Proportion of the market affected*

[1174] Google response to the interim report and MIR consultation, page 4.
[1175] Ofcom analysis of data provided in response to our information requests and data from Synergy and IDC. Annual growth based on the compound annual growth rate between 2019 and 2022.
[1176] The market research we conducted is described in Section 2.
[1177] IDC, 2023. Worldwide Black Book: Live Edition, July (V2 2023) Forecast (published July 2023). Total IT spending includes the following IDC technology categories from the Black Book publication: Infrastructure, Application Development & Deployment, Applications, System Infrastructure Software, Managed Services, Support Services, Project Oriented Services and Devices.

11.29 Our analysis suggests that the barriers to multi-cloud and switching we have identified will likely affect a large portion of the market. Around half of respondents in our customer survey only use a single cloud provider, and customers who multi-cloud rarely run highly integrated solutions across different clouds. In the case of switching, our market research found that 'difficulty and expense of switching providers' was the top concern about the cloud infrastructure market – being cited by more than half of respondents (59%). Concerns about egress fees and interoperability also ranked highly (55% and 52% respectively). Our evidence indicates that more mature and larger organisations, which are likely to make up most of the revenues of cloud providers, may be particularly affected by technical barriers. This may be because these customers are more likely to have large numbers of applications and/or use various proprietary services offered by their cloud providers, which add to the complexity of switching cloud provider. Evidence gathered from the hyperscalers is consistent with this.

11.30 Furthermore, in the UK, the hyperscalers collectively account for around 70% to 80% of total UK revenues generated from IaaS and PaaS.[1178] Of that figure, we estimate AWS and Microsoft Azure accounted for around 70% to 80% in 2022, with Google at around 5% to 10%; with both ratios growing since 2019.

11.31 These market shares indicate that by far the majority of customers of public cloud infrastructure services will use hyperscaler services to a greater or lesser extent. Therefore, features of the market that are present in the hyperscalers' service offerings have the potential to affect the vast majority of cloud users.

11.32 Apart from direct customers of public cloud infrastructure services, another heavy user group are independent software vendors (ISVs). ISVs are suppliers of cloud services that do not usually own the underlying infrastructure. In such cases, ISVs may rely on cloud providers as suppliers and as distributors (for example via a hyperscaler's marketplace). This means that the features we have identified also affect ISVs and their customers.

*Persistence of those features*

11.33 Our view, based on the evidence gathered in our market study, is that the features of the market we have identified may have a sustained long-term impact in the market. Where competitive constraints from new entrants have changed other behaviours (for example, lowering ingress fees to zero),[1179] incentives to address our areas of concern remain low.

11.34 The market is relatively concentrated and the market share of the leading providers AWS and Microsoft is consistently high. AWS has earned persistently high profits above our estimate of the weighted average cost of capital (WACC) and Microsoft Azure return on capital employed has increased since 2018 and is now above our estimate of WACC.[1180] We consider this evidence indicates it is likely that these features will persist in the future.

11.35 There are EU initiatives that could mitigate some of the concerns we have identified, specifically the EU Data Act and the Digital Markets Act. These changes do not apply directly to the UK market. However, we have not been able to observe their efficacy or influence in

[1178] See Annex 1.
[1179] Cloud providers have an incentive to make it easy and cheap to move data into their cloud, because revenues from the storage, compute and other services provided relies on this data being in the cloud. Indeed, AWS charged for ingress until Microsoft entered the cloud market.
[1180] For more detail, see Annex 2.

practice on market behaviours, and ultimately there is no evidence that they would change the dynamics of the UK market.

*Conclusion on the first criterion: scale of the suspected problem*

11.36 We consider that the scale of the suspected problems in relation to public cloud infrastructure services has a high likelihood of adverse effects on competition which affect a large proportion of a significant market.

## Second criterion: availability of appropriate remedies through an MIR

11.37 At this stage, we consider there may, in principle, be a number of appropriate remedies to the competition concerns and resulting detrimental effects of the above features. In this market study, our role is to assess the availability of potential remedies but not to assess their appropriateness in any detail. This would be the role of the CMA in its investigation if it finds an adverse effect on competition in relation to public cloud infrastructure services in the UK.[1181] We received a range of responses on potential interventions, which we deal with in detail in this report.[1182] In summary, many stakeholders were generally supportive of one or more intervention options we cited.[1183] The hyperscalers and some other stakeholders were less supportive and were concerned about the unintended consequences of regulatory intervention or preferred an industry first approach.[1184]

11.38 A non-exhaustive list of potential remedies that a market investigation could consider includes:

*Egress fees*

a) Equalise egress fees with other charges (e.g. no higher than data transfer costs within a single provider's cloud).
b) Place a price control that restricts egress fees to 'at cost' charges.
c) Prevent providers from charging for data egress.

*Interoperability and portability*

d) Requirements for suppliers of cloud services to be more transparent about the interoperability of their cloud services.
e) Requirements for suppliers of cloud services to make their cloud services easier to interoperate with.

---

[1181] Without fettering the discretion of the CMA to identify alternative approaches or remedies we provide further detail of our assessment of each of the potential remedies set out here in Section 10 of our final report.
[1182] See in particular, Section 10.
[1183] BT Group response to the interim report and MIR consultation, page 1; Name Withheld 1 response to the MIR consultation, paragraph 3.1; Oracle response to the MIR consultation, page 3; Name Withheld 2 response to the MIR consultation, paragraph 3.1; Federation of Communication Services response to the interim report, paragraph 5.1 and 8.1; Genera8 response to the interim report, paragraph 5.1 and 8.1; Sustainable Digital Infrastructure Alliance e.V response to the interim report, paragraph 8.1, 8.2 and 8.3; Vodafone response to interim report, page 3; Cloudflare response to the interim report, page 1-3; [✂] response to the interim report, paragraph [✂]; [✂] response to the interim report, paragraph [✂]; and [✂] response to the interim report, paragraph [✂].
[1184] UKCTA response to the MIR consultation, pages 1-2; and [✂] response to the interim report, page [✂]. AWS response to the interim report, pages 6-14. Google response to the interim report and the MIR consultation, pages 2-4; Microsoft response to the interim report and the MIR consultation, pages 9-10; and University of East Anglia response to the interim report, page 3-5.

f) Requirements for suppliers of cloud services which aim to increase the degree of standardisation.
g) Requirements for cloud providers to interconnect their data centres.

*Committed spend discounts*

h) Prohibit or restrict discount structures which create an incentive to concentrate spending with a single provider and risk distorting competition.

11.39 In carrying out a market investigation process, the CMA has wide-ranging powers to accept undertakings or impose an Order, as well as to make recommendations. As highlighted by the examples above, we consider that there are a number of potential appropriate remedies within the scope of such powers.

*Conclusion on the second criterion: availability of appropriate remedies through an MIR*

11.40 We consider that appropriate remedies are likely to be available. As with all interventions of this potential scale and significance, the design and any ongoing involvement by the regulatory authority would need to be considered carefully prior to implementation.

## Third criterion: the availability of undertakings in lieu of a reference

11.41 Ofcom has the power under section 154 EA02 to accept UILs instead of making an MIR. Before doing so, Ofcom is obliged, pursuant to s.154(3) EA02, to: "have regard to the need to achieve as comprehensive a solution as is reasonable and practicable to the adverse effect on competition concerned and any detrimental effects on customers so far as resulting from the adverse effect on competition".

11.42 As the CMA's guidance notes, such UILs are "unlikely to be common", but "where an adverse effect on competition arises from the conduct of very few firms there may be more scope for accepting undertakings in lieu" than "when the adverse effects on competition arise from market features involving several firms or industry wide practices".

11.43 We have not received any offers of UILs from the relevant stakeholders.

## Fourth criterion: alternative powers available to Ofcom or the CMA

11.44 Finally, we have considered whether alternative powers are available to us, or others, and if so, whether it would be more appropriate to use those to address the features we have identified.

11.45 For the reasons set out in this report, our view is that a MIR is the most appropriate tool to address our concerns in relation to the features that we identify above. This is because where an adverse effect on competition is found, action can be taken to remedy competition issues and their harmful effects on a forward-looking basis, rather than seeking redress or imposing a sanction for past conduct (for which enforcement action under the CMA's other powers might be more appropriate). We believe that an MIR will allow the CMA to evaluate the range of the factors which give rise to competition concerns in relation to public cloud infrastructure services in the UK in a timely manner.

11.46 We have considered Ofcom's and the CMA's powers in relation to competition law prohibitions on anticompetitive agreements or abuse of a dominant position and in relation to consumer law, before considering the powers available to other regulators. We have not identified any grounds to suggest that it would be more appropriate or effective to address one or more of the features we have identified or their effects using competition or consumer powers.

11.47 Finally, the Digital Markets, Competition and Consumer Bill will establish a new pro-competition regime for digital markets. The Government has published draft legislation, which is currently undergoing Parliamentary scrutiny.[1185] As set out in Section 2, the Bill will empower the CMA to designate firms providing digital activities with strategic market status (SMS) and apply binding conduct requirements, to manage the effects of market power. The CMA may also apply pro-competitive interventions, to tackle the root causes of their market power. The Bill is not yet law and, as noted, its provisions are currently undergoing scrutiny in Parliament. Given this, in our view, it would not be appropriate to rely on the Bill to address the competition concerns that we have identified. Further consideration of the application of the Bill will be a matter for the CMA to consider in the course of the market investigation.

*Conclusion on fourth criterion: alternative powers*

11.48 We do not currently consider that alternative powers, or another regulator, could more appropriately address the concerns we have identified.

## Decision on an MIR

11.49 For the reasons set out above, we consider that it is appropriate to exercise our discretion to make an MIR in relation to cloud infrastructure services. Therefore, we have decided to make an "ordinary" MIR within the meaning of section 131(6) EA02 in respect of the supply of public cloud infrastructure services in the UK.

11.50 Although we have identified particular features, our MIR does not confine the CMA's market investigation to these features, and it will be for the CMA to decide whether to investigate other issues during the market investigation.[1186]

[1185] The relevant parliamentary materials, including drafts of the DMCC Bill are available on the UK Parliament website: Digital Markets, Competition and Consumers Bill [accessed 14 September 2023].
[1186] See the terms of reference set out in Annex 6.