
Cloud services market study

Interim report

Non-confidential version – redacted for publication [X]

CONSULTATION:

Publication date: 5 April 2023

Closing date for responses: 17 May 2023

Contents

Section

| | |
|---|-----|
| 1. Overview | 1 |
| 2. Introduction | 11 |
| 3. Market context | 16 |
| 4. How competition works | 42 |
| 5. Potential barriers to effective competition | 79 |
| 6. Potential competition concerns | 178 |
| 7. Microsoft's licensing practices | 192 |
| 8. Overview of potential interventions | 200 |
| 9. Next steps | 211 |
| 10. Our decision to consult on a market investigation reference | 213 |

Annex

| | |
|-------------------------------------|-----|
| A1. Responding to this consultation | 215 |
| A2. Ofcom's consultation principles | 218 |
| A3. Consultation coversheet | 219 |
| A4. Consultation questions | 220 |

1. Overview

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 is carrying 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 consultation sets out our interim findings and provisional view on recommendations. We are seeking input on our assessment of competition in the market and proposal for further investigation by the Competition and Markets Authority (CMA) following the study.

Our interim findings – 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 £4.5 billion to £5.0 billion in 2021.

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 60% to 70% in 2021.² 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.

At this interim stage of our study we have found some evidence of active competition in cloud infrastructure, especially where providers are competing to attract new customers who are moving to the cloud for the first time. There are indications that this is contributing to benefits for customers, including product innovation, discounts and a wide choice of software services from ISVs.

However, our provisional 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:

¹ We have published a [market study notice](#) in accordance with section 130A of the Enterprise Act 2002 as amended and applied by section 370 of the Communications Act 2003.

² We used a combination of data sources to estimate market shares and present our estimates in ranges for confidentiality reasons. See Annex 5 for more details.

- **egress fees** are the charges that customers pay to transfer their data out of a cloud and the hyperscalers set them significantly higher than most other providers. The cost of egress fees can discourage customers from using services from more than one cloud provider and in some cases can make it more costly to switch.
- **technical restrictions on interoperability** are imposed by the leading firms that prevent some of their services working effectively with services from other providers. This means customers need to put additional effort into reconfiguring their data and applications to work on different clouds.
- **committed spend discounts** can benefit customers by reducing their costs, but the way these discounts are structured can incentivise customers to use a single hyperscaler for all or most of their cloud needs. This can make it less attractive to use rival providers as new needs emerge.

As a result, we are concerned that a significant number of customers, especially those with more complex requirements, may face material barriers to switching and multi-cloud. This could leave some customers 'locked in' to one of the leading providers. We expect this will be true of an increasing number of customers as the market matures. We are most concerned in relation to AWS and Microsoft, given their market position and the fact they display some form of all the above behaviours that limit competition.

There are indications these market features are causing harm today, with some existing cloud customers paying more or settling for lower quality services, which in turn can lead to negative impacts for UK consumers. We see evidence of customers already in the cloud facing significant price increases when they come to renew their contracts. We have also heard concerns from some customers about their ability to switch and use multiple providers, which limits their access to the best quality products. 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.

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. In this scenario we are concerned that the threat of customers switching away from the market leaders will reduce, further dampening competition for new and existing customers. 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 think it is appropriate to assess the case for taking action now, both to improve outcomes today and reduce risks in the future. We are consulting on a proposal to refer the cloud infrastructure market to the Competition and Markets Authority (CMA) to carry out a market investigation. This would allow the CMA to further examine if there are interventions that could address the adverse impact of the barriers we have identified and improve how the market works for customers.

We have also heard concerns about Microsoft's licensing practices in relation to some of the enterprise software services it provides to business customers. These allege that Microsoft sells certain software services in a way that makes it less attractive to use them on rival cloud infrastructure compared to its own called Azure. Microsoft disputes the veracity of the concerns. While the concerns raised stem from a market outside the scope of our study, it is possible they could dampen competition in cloud infrastructure. Ofcom and the CMA will consider the most appropriate way forward on these issues.

We welcome views on the provisional findings in our interim report. Alongside this document we are consulting on a proposal to refer the cloud infrastructure market to the CMA for further investigation. We will complete our study and publish the final report by 5th October 2023.

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

- 1.1 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.
- 1.2 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.
- 1.3 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.
- 1.4 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.
- 1.5 Our work complements ongoing initiatives on cloud computing by other UK and international regulators, including other market wide studies.

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 60% to 70% 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 above our estimate of the weighted average cost of capital (WACC) since at least 2015. 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 now 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. Although its financial performance is improving, Google's cloud division is yet to make a profit.
- 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. Collectively they account for a significant share 30% to 40% of UK PaaS revenues, but no individual company has more than 0% to 5%.

Competition is 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 25% to 30% 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 very 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 currently 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.³ 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 that the behaviour of market leaders could increase 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 practices that appear to raise barriers to effective competition by making it more difficult for customers to switch and multi-cloud than might otherwise be the case. The practices we are most concerned about are the charging of egress fees, restrictions on interoperability and the structure of committed spend discounts. These are likely to most severely affect customers with more complex requirements. We are most concerned about their impact in relation to AWS and Microsoft, given they have a strong position in the market and undertake all the practices we have identified in some form.
- 1.18 Our provisional view is that these practices, either alone or acting in combination, can limit the ability of customers to switch provider or adopt a multi-cloud strategy. In each case, further investigation is needed to establish the extent to which they represent a significant barrier to competition.

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

- 1.19 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 cloud. These charges are known as egress fees. Egress fees can be particularly significant for cloud applications hosted on different clouds that need data to be moved regularly between them. 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 want to switch their

³ 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).

primary cloud provider, particularly where this involves a large one-off transfer of data between clouds.

- 1.20 The hyperscalers charge a similar level of egress fees, which are 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 the hyperscalers appear to set the level of their egress fees above the underlying cost of transferring data.
- 1.21 Egress fees are a key concern for existing customers because they may 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.

The interoperability of some cloud services is being restricted in ways that limit the ability of customers to combine products from different providers

- 1.22 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 interoperability works well, it can unlock significant benefits for customers. However, a lack of interoperability 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 either combine different services across cloud providers or change their primary provider. It can also inhibit customers from adding a secondary provider to serve specific uses cases as their requirements evolve.
- 1.23 While a degree of complexity is inevitable given the nature of cloud services, we have seen evidence indicating that technical barriers are more significant than they need to be. For example, we have seen evidence that AWS and Microsoft may limit the interoperability of some of their services by not openly sharing important technical information needed to interoperate with their service. This means they can only be used in their respective clouds, which prevents customers from using them in combination with products from rival providers.
- 1.24 We have also heard that providers may be unnecessarily limiting the interoperability of their cloud services by using proprietary or 'closed' cloud technologies. In some cases, providers adopt open standards or application programming interfaces (APIs) initially, but then introduce modifications over time that have the effect of integrating the service more tightly to their ecosystem. These limitations are not always apparent to customers when they make the initial purchase.
- 1.25 In both instances, some customers are prevented from switching or building their preferred cloud architecture, where they can mix and match the cloud services that most

closely meets their needs. This is reflected in our customer research, which found that around half of customers (52%) are concerned about a lack of interoperability.

The structure of committed spend discounts can encourage some customers to use a single hyperscaler for 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 could act as a barrier to entry and expansion by encouraging mainly 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 In practice when a customer has new workloads to move into the cloud, the prospect of losing their existing committed spend discount could make it less attractive to use a rival provider.⁴ We think this is a particular concern where customers face barriers to switching their existing cloud use. Ultimately this could restrict competition by hampering the ability of rival providers to compete effectively for any new workloads as they emerge.

Some customers face material barriers to switching and are experiencing harm today

- 1.29 Where there is active competition for new customers, in particular 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 the barriers we have identified are then strong enough to result in a material number of customers having a limited ability to switch or use multiple providers, to the extent that they consider themselves to be 'locked-in' to their current provider. Our research found that lock-in is already a real concern for customers and is likely to be a greater risk for those businesses with more complex cloud requirements.
- 1.30 Where customers face material barriers to switching and multi-cloud, they could find themselves in a weak bargaining position with their existing provider, either when they renew their contract or purchase additional cloud services as their requirements evolve. This could weaken competitive constraints on the market leaders as customers cannot

⁴ A workload is a specific application, service, capability or a specific amount of work that can be run on a cloud resource.

credibly threaten to switch away, with these customers paying more than they should or not purchasing the best quality services from a rival provider. Evidence indicates that price rises are relatively common when a customer re-negotiates terms with their provider and can be significant (20% on average). In principle, customers can protect themselves from these future risks when they initially contract with their provider, but their ability to do so may be limited due to factors such as the difficulty of forecasting their future demand for cloud services.

- 1.31 At the same time, we estimate that AWS and Microsoft Azure earn profits above the estimated cost of capital, which indicates that they can maintain high prices for certain customers. Our analysis indicates that AWS's profitability has been consistently high, with returns above the weighted average cost of capital (WACC) since at least 2015. We estimate that Microsoft's Azure returns have increased in recent years and are now above our estimate of WACC.
- 1.32 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. Limits on entry and innovation can directly affect the quality and range of choice UK consumers have of online services powered by the cloud.

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

- 1.33 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.
- 1.34 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 start to 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 ecosystems.
- 1.35 Today we see some evidence that the market leaders have an incentive to compete to win new customers and to a lesser extent for additional workloads to retain those customers within their ecosystems. A weaker competitive constraint from rivals 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.
- 1.36 In a more concentrated market, the leaders may 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 rival ISVs, for example by acting in

ways that benefit their own competing products. In turn this could impact the choice, quality and prices that ISVs are able to offer to their customers.

In response we propose to refer the cloud infrastructure market to the CMA for an in-depth investigation

- 1.37 Our assessment at this stage of the study is that, while there are some positive signs of competition at present, there are also clear indications that the cloud infrastructure market may not be working well. We have provisionally identified features of the market that we suspect 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. We think there are credible interventions that could address our concerns and it is important to assess the case for taking action now, both to improve outcomes today and reduce risks in the future.
- 1.38 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. Our provisional view is that these features may be having an adverse effect on competition.
- 1.39 On this basis we are consulting on a proposal to refer the market for cloud infrastructure services to the CMA to carry out a market investigation.⁵ In reaching this proposal we have assessed our concerns in line with CMA guidance on market investigation. Our initial assessment is that that legal threshold is met and a market investigation reference is an appropriate response to the concerns we have identified.
- 1.40 Making a market investigation reference would be a significant step for us to take. Our proposal reflects the importance of cloud computing to UK consumers and businesses, the significant concerns we have about the cloud infrastructure market and our view that the CMA are best placed to undertake any further investigation. Should it find an adverse effect on competition, the CMA would decide whether to take such action as it considers reasonable and practicable to 'remedy, mitigate or prevent' any adverse effects on consumers. In response the CMA may impose a broad range of remedies. We have identified potential interventions that could be implemented, but it would ultimately be up to the CMA to assess their proportionality and effectiveness as part of any market investigation.
- 1.41 It is possible that undertakings may be offered following publication of this consultation that address the concerns we have identified. If this were the case, we would consider them.

⁵ Ofcom, 2023. [Public cloud infrastructure services, Consultation: Proposal to make a market investigation reference.](#)

We have also heard concerns about the potential for software licensing practices to impact competition in cloud infrastructure

- 1.42 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).
- 1.43 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 impacts on their ability to compete for customers. Microsoft disputes the veracity of the concerns.
- 1.44 It is possible that the alleged conduct could risk dampening competition in cloud infrastructure services. Reaching a view on that would require detailed examination of the concerns raised, which stem from Microsoft’s position in an adjacent market – the supply of enterprise software. Our market study is not the appropriate tool to undertake that assessment and we have not sought to make any findings in this document. Ofcom and the CMA will consider the most appropriate way forward on these issues.

Next steps

- 1.45 We look forward to receiving feedback from stakeholders on our interim report. In particular, we welcome responses to the questions set out in Annex 4.
- 1.46 Alongside this document we have also published a separate [consultation](#) on our proposal to make a market investigation reference into the supply of public cloud infrastructure services in the UK.
- 1.47 Both consultations are open for six weeks and close on 17 May 2023.
- 1.48 We will continue to gather evidence and engage with stakeholders during the second half of the study to inform our work. We intend to publish a final report setting out our findings and recommendations, including our decision on a market investigation reference, by no later than 5 October 2023.

2. Introduction

- 2.1 In this section, we provide some context to the market study, summarise the evidence gathering we have conducted so far and explain the purpose and structure of this interim report.

Context

- 2.2 We launched a market study into cloud services on 6 October 2022 using our powers under the Enterprise Act 2002. We did this to better understand the market for cloud services in the United Kingdom (UK) and determine whether it works well for consumers.⁶
- 2.3 ‘Cloud computing’ is 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. Compared to traditional information technology (IT), cloud computing offers flexibility which enables customers to quickly scale up or down the computing resources that support their business. ‘Cloud services’ are all services involved in the provision of cloud computing.
- 2.4 Cloud services are a fundamental input across all sectors of an increasingly digitised economy including the telecoms and broadcasting sectors which are central to our remit. They play a critical role in the delivery of a range of communications services to consumers over the internet.
- 2.5 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
- 2.6 We launched a market study, using our powers to exercise the Competition and Markets Authority’s (CMA’s) functions under Part 4 of the Enterprise Act 2002. Ofcom has concurrent functions with the CMA. Section 370(3A)(b) of the Communications Act 2003 (‘CA03’), when read together with section 130A of the Enterprise Act 2002 (‘EA02’), gives Ofcom the power to undertake a Market Study. This Market Study considers the extent to which a matter in relation to commercial activities connected with communication matters has, or may have, effects adverse to the interests of consumers.

⁶ Ofcom, 2022. [Cloud services market study notice](#).

⁷ In September 2022, we published a document setting out our approach to competition and consumer issues in internet-based communications markets - see Ofcom, 2022. [Digital markets in the communications sector](#). We set out a programme of work which includes four targeted projects that are of particular importance to UK consumers and the functioning of communications markets. They are net neutrality, competition in digital content gateways, online personal communications services and this market study on cloud services.

2.7 Cloud services are of interest to several jurisdictions around the world given their role in the global digital economy, including existing and emerging legislative developments within the European Union (EU). As part of our study to date, Ofcom has been in dialogue with the French, Dutch, and Japanese competition authorities on their respective cloud market reviews.⁸ In addition, we note that the United States (US) Federal Trade Commission (FTC) has recently issued a request for information seeking public comment on the business practices of cloud computing providers, with a particular interest in the healthcare, finance, transportation, e-commerce and defence sectors.⁹

Scope

- 2.8 We launched the study with a call for inputs (CFI) inviting views from stakeholders on a range of issues including its scope.¹⁰
- 2.9 We proposed to focus on ‘public cloud’ (where cloud services offer access to a shared pool of computing resources, rather than being reserved for a single customer), as these represent the biggest transformation to the way that businesses buy computing resources.
- 2.10 We also proposed concentrating on ‘cloud infrastructure services’ which include services that provide access to raw computing resources, i.e. basic compute, storage and networking (often referred to as infrastructure as a service, or 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).
- 2.11 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 may pose a risk to effective competition. Competition concerns in this part of cloud could have far-reaching effects for the wider economy, particularly where vertically integrated cloud providers act both as a supplier of cloud infrastructure services and a competitor to third-party SaaS providers.
- 2.12 Respondents were broadly supportive of our focus on public cloud and cloud infrastructure although some felt that the scope should be expanded to include issues of cloud security and resilience. We recognise that these are factors in customer decision making when choosing cloud services, deciding how to deploy them and selecting providers and we discuss this in Section 4. However, it is not feasible to examine all aspects of cloud services in this study and therefore we have retained our focus on competition. Although we expect our findings to complement and support wider policy work on security resilience.

Evidence-gathering

2.13 During the first few months of the study, we focused on gathering information as follows.

⁸ See paragraphs 2.19 to 2.24 of our CFI.

⁹ [FTC Seeks Comment on Business Practices of Cloud Computing Providers that Could Impact Competition and Data Security | Federal Trade Commission](#) [accessed 23 March 2023].

¹⁰ Ofcom, 2022. [Cloud services market study, call for inputs](#).

Call for inputs

- 2.14 As mentioned above, we began by publishing a CFI seeking views on the proposed scope of the study and on:
- our initial characterisation of the market;
 - our proposed approach for considering the dynamics in cloud infrastructure services competition and to examine cloud ecosystem competition;
 - any concerns regarding any conduct or activities of any provider(s) that may adversely affect market dynamics now or in the future; and
 - any remedies that we should investigate further to mitigate some of the potential risks or concerns with the market.
- 2.15 We received 11 responses to our CFI and published non-confidential responses on our website.¹¹

Market research

- 2.16 To help to us better understand the customer perspective, we commissioned some market research which we have published alongside this interim report.¹² This included both qualitative and quantitative phases to produce robust insights and 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 questionnaire

- 2.17 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.

Telecoms and broadcasting

- 2.18 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 statutory powers to request information from them.¹³

Suppliers in the cloud services value chain

- 2.19 Turning to suppliers, we have had discussions with the hyperscalers¹⁴ 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

¹¹ Non-confidential responses to our CFI are published on our [website](#).

¹² Context Consulting, 2023. [Cloud Services Market Research, Summary of Findings](#) and [Data Tables](#).

¹³ Section 174 of the Enterprise Act 2002.

¹⁴ Amazon Web Services, Microsoft and Google.

independent software vendors (ISVs) and providers of professional services such as resellers, consultants and managed service providers.

Market analysts and public information

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

Other regulators

2.21 In our CFI we set out wider UK policy interests concerning cloud including work then being undertaken by other regulators.¹⁵

2.22 Since launch we have continued to consult with the Competition and Markets Authority (CMA) as we have progressed our work and will continue to do so during the second half of the study¹⁶. We have also engaged with other regulators including the Information Commissioner's Office (ICO), who responded to our CFI,¹⁷ and the Financial Conduct Authority (FCA),¹⁸ both bilaterally and through the Digital Regulation Cooperation Forum (DCRF).¹⁹

2.23 We are grateful to all parties who have engaged with the study and helped us make substantial progress over the last six months.

2.24 We will continue to engage with stakeholders and, where appropriate, conduct further evidence gathering in the coming weeks.

Purpose of the interim report

2.25 Published half-way through our market study, the purpose of this interim report is to provide an update on our approach and our progress, to indicate the direction of travel our analysis is taking in relation both to concerns and potential interventions to address them, and to test these initial findings with stakeholders.

2.26 We have structured our interim report as follows:

- Section 3 provides context about the cloud services market, providing an overview of the service and deployment models, and outlines the main players in the market.
- Section 4 details how competition works in the sector and considers key market outcomes (such as market shares and profitability) and assesses the extent to which they may be indicative that competition may not be as effective as it could be.
- Section 5 covers the various potential barriers to competition in the sector, including barriers to switching and multi-cloud, barriers to entry and expansion, and hyperscalers' relationship with ISVs.

¹⁵ Ofcom, 2022. CFI, para 2.14 to 2.18.

¹⁶ [About us - Competition and Markets Authority - GOV.UK \(www.gov.uk\)](https://www.gov.uk/about-us) [accessed 16 March 2023].

¹⁷ ICO, 2022. [The Information Commissioner's response to Ofcom's cloud services market study call for inputs](https://www.ico.gov.uk/about-ico/news-and-events/2022/03/22/the-information-commissioner-s-response-to-ofcom-s-cloud-services-market-study-call-for-inputs).

¹⁸ <https://www.fca.org.uk/> [accessed 16 March 2023].

¹⁹ <https://www.gov.uk/government/collections/the-digital-regulation-cooperation-forum> [accessed 16 March 2023].

- Section 6 summarises our findings on the current state of competition in the sector, highlighting our concerns about where it is weak and, if left unchecked, how competition in the market could deteriorate.
 - Section 7 summarises submissions we have received regarding Microsoft’s software licensing practices and articulates the relevance of submissions for competition in cloud infrastructure.
 - Section 8 identifies features of the market where we believe there are strong arguments for intervention and a summary of what that might look like.
 - Section 9 sets out our provisional view of market study outcomes, and a high-level plan of work for the next six months until final report.
 - Section 10 sets out the rationale for our decision to consult on a proposal to make a market investigation reference to the CMA, in relation to the supply of public cloud infrastructure services in the UK.
- 2.27 Details of how to respond to this six-week consultation on our interim report are set out in Annexes 1 to 4. The deadline for responding to this consultation is 17 May 2023.
- 2.28 We will take views into account before reaching our conclusions and recommendations. A final report will be published by 5 October 2023.
- 2.29 Alongside this interim report we are also publishing a consultation document on our proposal to make a market investigation reference to the CMA.²⁰

²⁰ Ofcom, 2023. [Public cloud infrastructure services, Consultation: Proposal to make a market investigation reference](#).

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 and broadcasting 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)²¹ 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.”²² 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”.²³

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.

²¹ NIST is a non-regulatory agency within the US Department of Commerce. The NIST definition of cloud computing is widely adopted.

²² NIST, 2011. [The NIST Definition of Cloud Computing](#) [accessed 15 February 2023].

²³ DCMS, 2022. Policy paper. [Data storage and processing infrastructure security and resilience - call for views](#) [accessed 16 February 2023].

- 3.6 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,²⁴ which provide a full 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 full range of cloud services, while others are more specialised. We discuss suppliers in more detail later in this section.
- 3.7 Worldwide end user spending on public cloud services is forecast to grow 20.7% to total £477.7 billion in 2023, up from £395.7 billion in 2022.²⁵ In comparison, worldwide IT spending is projected to total £3.6 trillion in 2023, an increase of 2.4% from 2022.²⁶ Based on the available data, it appears that public cloud spending continues to accelerate. In addition, public cloud spend makes up a significant proportion of IT budgets, especially among large organisations.

Importance of cloud services in the UK

- 3.8 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.
- 3.9 Cloud is also a cornerstone of recent technological innovations. From data science to artificial intelligence (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.
- 3.10 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,²⁷ increase their innovation potential,²⁸ enhance their quality of service, and achieve baseline security and resilience. It also offers access to relevant data from any device, anytime and anywhere.

²⁴ 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.

²⁵ Gartner Press Release, [Gartner Forecasts Worldwide Public Cloud End-User Spending to Reach Nearly \\$600 Billion in 2023](#), 31 October 2021 [accessed 29 March 2023]. These figures have been converted into pound sterling using ONS exchange rate data ([ONS Average Sterling exchange rate: US Dollar](#)).

²⁶ Gartner Press Release, [Gartner Forecasts Worldwide IT Spending to Grow 2.4% in 2023](#), 18 January 2023 [accessed 29 March 2023].

²⁷ 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.

²⁸ 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 <https://aws.amazon.com/about-aws/sustainability/> [accessed 16 February 2023].

- 3.11 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. Public cloud computing is completely reliant on internet connection, so if the connection is interrupted, data cannot be accessed.
- 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.²⁹ For example:
- Our customer research found that 82% of respondents had increased their spend on cloud in recent years, with 26% having greatly increased their budget.³⁰
 - Flexera research in 2023 with (mainly larger) European organisations found that more than half (62%) of respondents stated they were using cloud heavily.³¹
 - 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 2025.³²
- 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:
- Our customer research found that 79% of respondents expect to spend more on cloud in the next 18 months.³³
 - A 2022 paper by [redacted] 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.³⁴
 - 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.³⁵
- 3.14 While data on the extent of migration to cloud in the UK is more limited, we expect that most larger UK businesses and organisations will have already started migrating workloads to cloud to some degree. Many existing cloud customers are likely to migrate more of their

²⁹ 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, [Public Sector Cloud Strategy Report](#) (2022).

³⁰ 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.

³¹ Flexera, 2023. [2023 State of the Cloud Report](#), page 75 [accessed 31 March 2023]. It is worth noting that a large share (62%) of these European respondents were from the UK.

³² Gartner Press Release, [Gartner Says More Than Half of Enterprise IT Spending in Key Market Segments Will Shift to the Cloud by 2025](#), 9 February 2022 [accessed 29 March 2023].

³³ Context Consulting research report, Slide 42.

³⁴ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], [redacted].

³⁵ Goldman Sachs Equity Research, 12 July 2022. IT Spending Survey.

workloads in the next few years, though not all workloads will be appropriate for cloud migration.

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:
- Computing resources – these include hardware (servers and network equipment) and software (applications) which are used to process workloads³⁶ and store data.
 - 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.
 - 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.³⁷ Some years later Microsoft followed suit with Microsoft Azure in 2010 and Google with Google Cloud in 2011.³⁸

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 service models have been recognised by the NIST.³⁹ 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 control⁴⁰ the customer has over the management and maintenance of the computing resources. 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.

³⁶ A workload is a specific application, service, capability or a specific amount of work that can be run on a cloud resource.

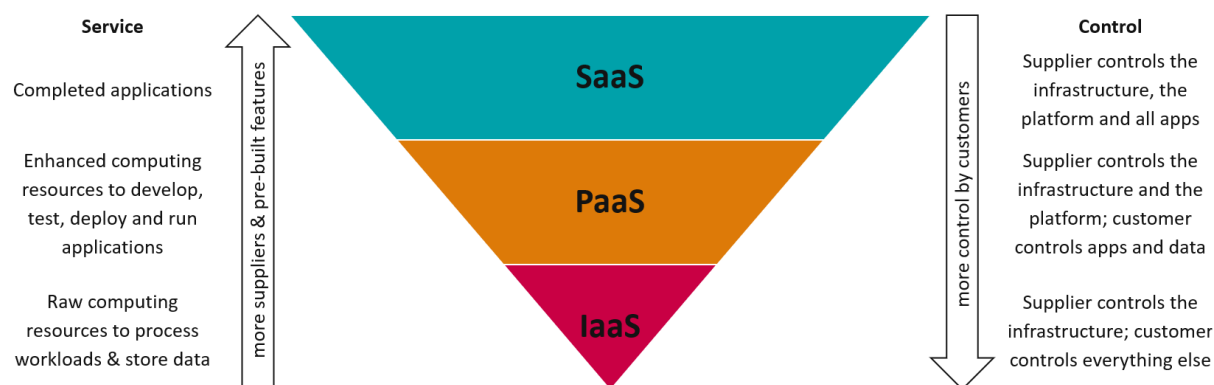
³⁷ AWS website. [About AWS](#) [accessed 20 February 2023].

³⁸ AWS, Microsoft and Google offered some cloud services in beta version before the official launch of their clouds.

³⁹ [The NIST Definition of Cloud Computing](#) [accessed 20 February 2023].

⁴⁰ 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.

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.⁴¹ 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.⁴² 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 2021, UK IaaS revenues were around [§<] [£2.5 to £3.0] billion, and grew by 25% to 30% per year between 2019 and 2021.⁴³

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,⁴⁴ machine learning and IoT (internet of things). The overall virtual environment and the underlying raw

⁴¹ 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.

⁴² AWS website. [Amazon EC2](#) [accessed 20 February 2023]; Microsoft Azure website. [Virtual Machines](#) [accessed 20 February 2023]; and Google Cloud website. [Compute Engine](#) [accessed 20 February 2023].

⁴³ 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 2021. See Section 4 on UK shares of supply for more detail.

⁴⁴ 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.

computing resources are typically owned and managed by the same cloud provider.⁴⁵ 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.⁴⁶ We estimate that in 2021, UK PaaS revenues were around [£1.5 to £2.0] billion, and grew by 35% to 40% per year between 2019 and 2021.⁴⁷

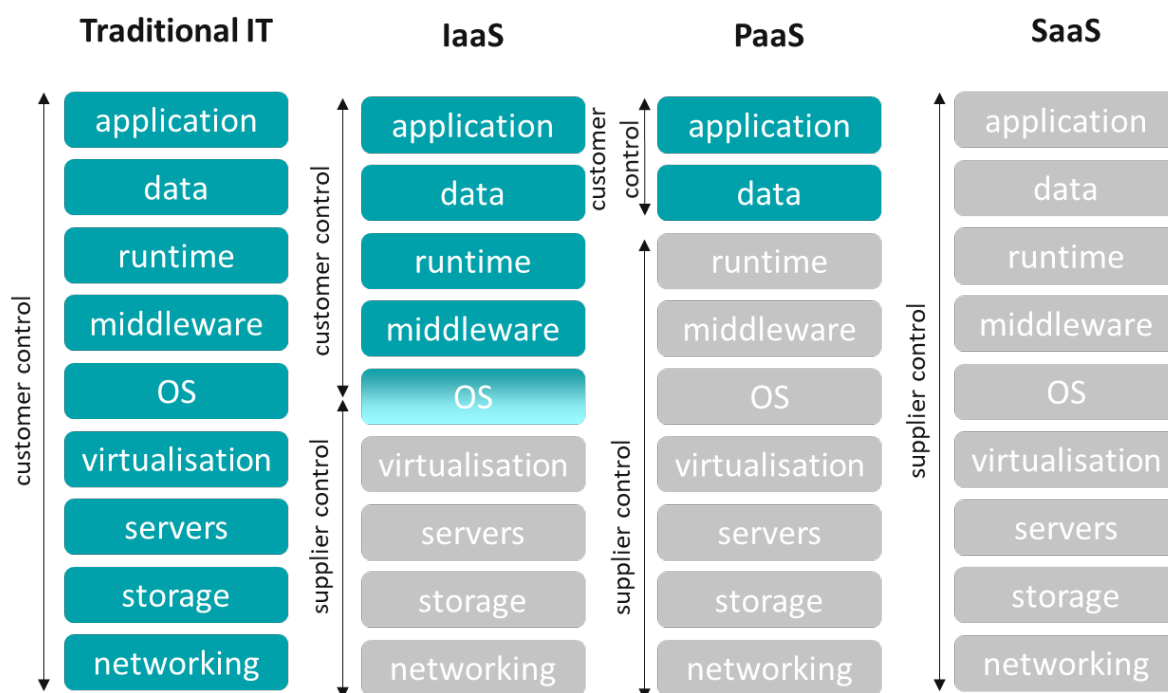
3.21 **Software as a service (SaaS)** are complete applications hosted in the cloud. These cloud applications can be offered by the cloud provider that owns the underlying raw computing resources or by an ISV. The provider of the SaaS service manages all hardware and software. In general, most modern 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.

⁴⁵ 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.

⁴⁶ AWS website. [AWS Elastic Beanstalk](#) [accessed 20 February 2023]; Microsoft Azure website. [DevOps solutions on Azure](#) [accessed 20 February 2023]; and Google Cloud website. [App Engine](#) [accessed 20 February 2023].

⁴⁷ 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 2021. 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 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.⁴⁸

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:

- 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

⁴⁸ Microsoft Azure website. [Virtual Machines](#) [accessed 23 March 2023]; IBM Storage as-a-Service [accessed 30 September 2022]; Red Hat website. [What is CaaS?](#) [accessed 23 March 2023]; Oracle Cloud Infrastructure (OCI) website. [What is DBaaS?](#) [accessed 23 March 2023]; and VMware website. [What is disaster recovery as a service \(DRaaS\)?](#) [accessed 23 March 2023].

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

- 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 legacy IT is not easily transferable to public cloud and for running latency-sensitive workloads close to the end-user. 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.
- 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 Multi-cloud is a cloud deployment model that involves the use of more than one public cloud provider by a single customer. The use of multiple public cloud providers can benefit customers by allowing them to access their preferred services, gain commercial bargaining power against their cloud providers and build for resilience.

3.25 In practice, as further detailed in Section 5, we understand that multi-cloud may take various forms and can be categorised as follows:⁴⁹

- 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 resiliency reasons, by duplicating some parts of their 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. different applications are 'siloed' on different public clouds). This multi-cloud architecture appears to be the most common amongst customers. It may be implemented to access 'best-of-breed' functionalities across clouds for specific

⁴⁹ 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 so far, and it is useful for our purpose of assessing the state of competition in cloud infrastructure services.

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 public cloud architecture by mixing and matching cloud services hosted on different public clouds. Unlike the cloud duplication and the siloed multi-cloud scenarios, in this case different cloud applications and data hosted on different public clouds are highly integrated. 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.

- 3.26 Our market research found that 52% of IaaS/PaaS users reported using more than one IaaS/PaaS provider.⁵⁰ This is likely to overstate the fraction of customers that use 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 Section 5 and Annex 7.

Customer preferences and behaviour

- 3.27 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.
- 3.28 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.⁵¹

Current and expected future cloud use

- 3.29 In our research the most frequently mentioned reason to adopt cloud computing is greater flexibility and agility.⁵² The second most important driver is improved security. These two

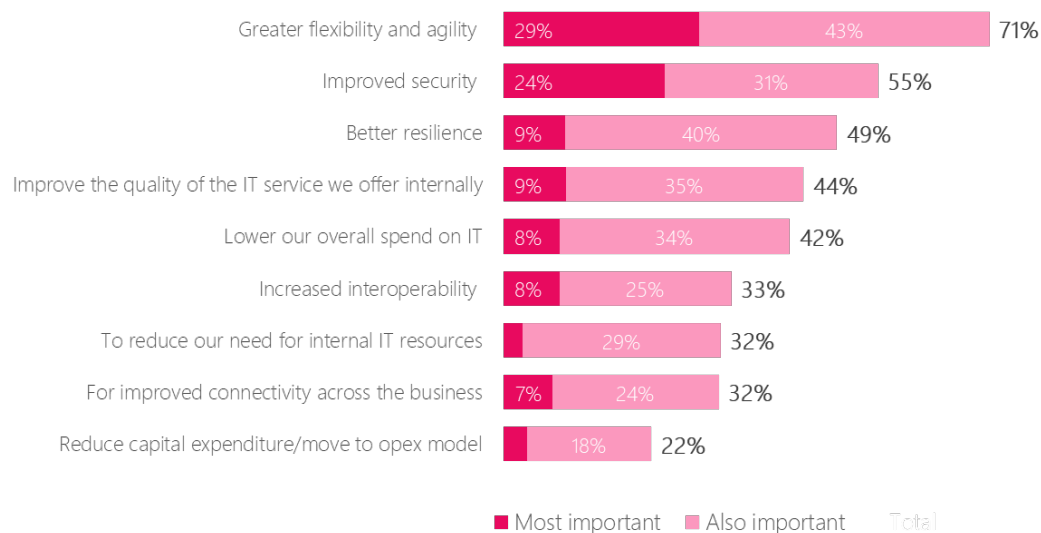
⁵⁰ Context Consulting research report, Slide 75.

⁵¹ We refer to the results of our research throughout our consultation document. 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.

⁵² Context Consulting research report, Slide 29.

drivers are also the two most frequently mentioned when respondents were asked to identify their single most important reason to adopt cloud computing. Financial reasons appear relatively less important: reduction of capital expenditure 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

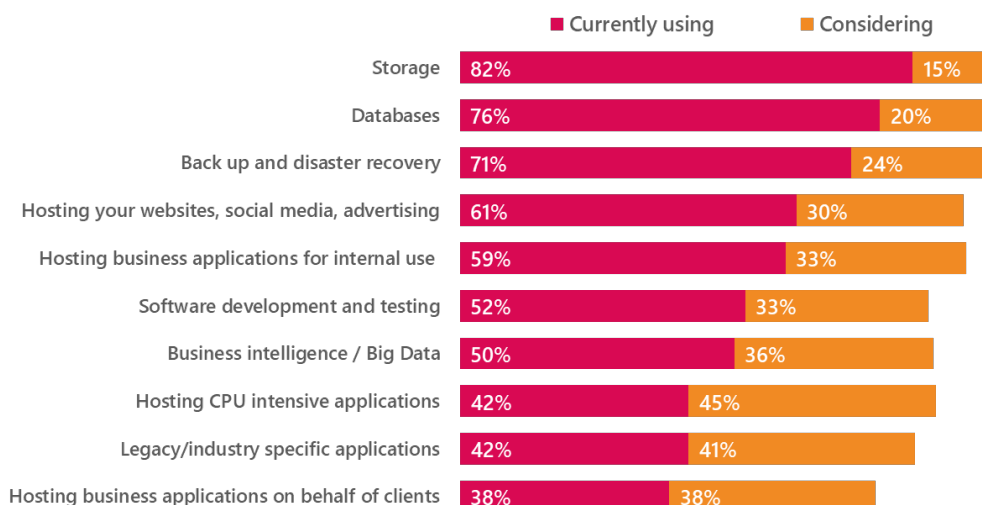


Source: Context Consulting research report, Slide 29.

- 3.30 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 a customer's business needs and the level of control they would like to have over their cloud service.
- 3.31 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, 55% - PaaS.⁵³
- 3.32 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%).

⁵³ 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.

- 3.33 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 3 use cases or fewer, 31% have 4-5, 31% have 6-7, and 22% have 8-10. In general, larger organisations, as well as IT & Tech companies, tend to have more cloud use cases.⁵⁴
- 3.34 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.⁵⁵

⁵⁴ Context Consulting research report, Slide 37.

⁵⁵ “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.

3.35 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 into cloud as top priorities for companies.⁵⁶

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

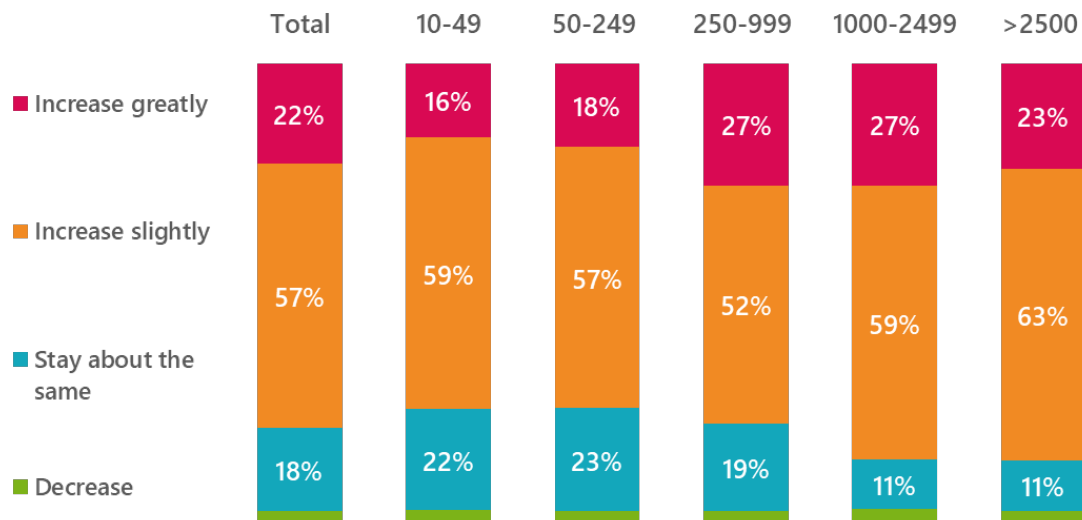
3.36 For 82% of respondents spend on cloud services has increased in recent years (for 26% it increased greatly, and for 57% - slightly). 16% of respondents reported that their spend stayed about the same, and for about 1% it decreased. The increase in spend is consistent across different respondent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

3.37 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%).

3.38 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).

⁵⁶ Flexera, 2023. 2023 State of the Cloud Report, page 29.

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

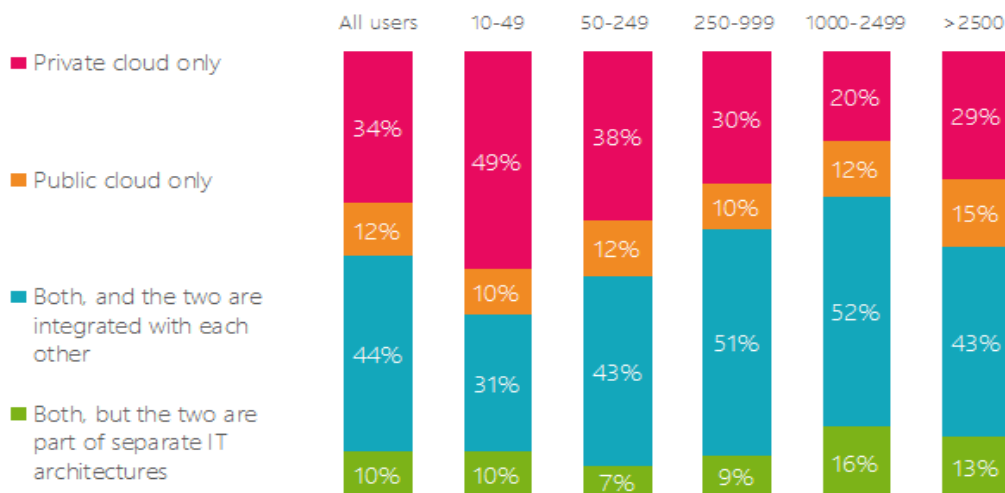


Source: Context Consulting research report, Slide 42.

Use of private, public and hybrid cloud

3.39 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.

3.40 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 3 or more cloud providers tend to use integrated public and private cloud more than others.

- 3.41 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.
- 3.42 In contrast, the Flexera 2023 State of the Cloud report⁵⁷ 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.

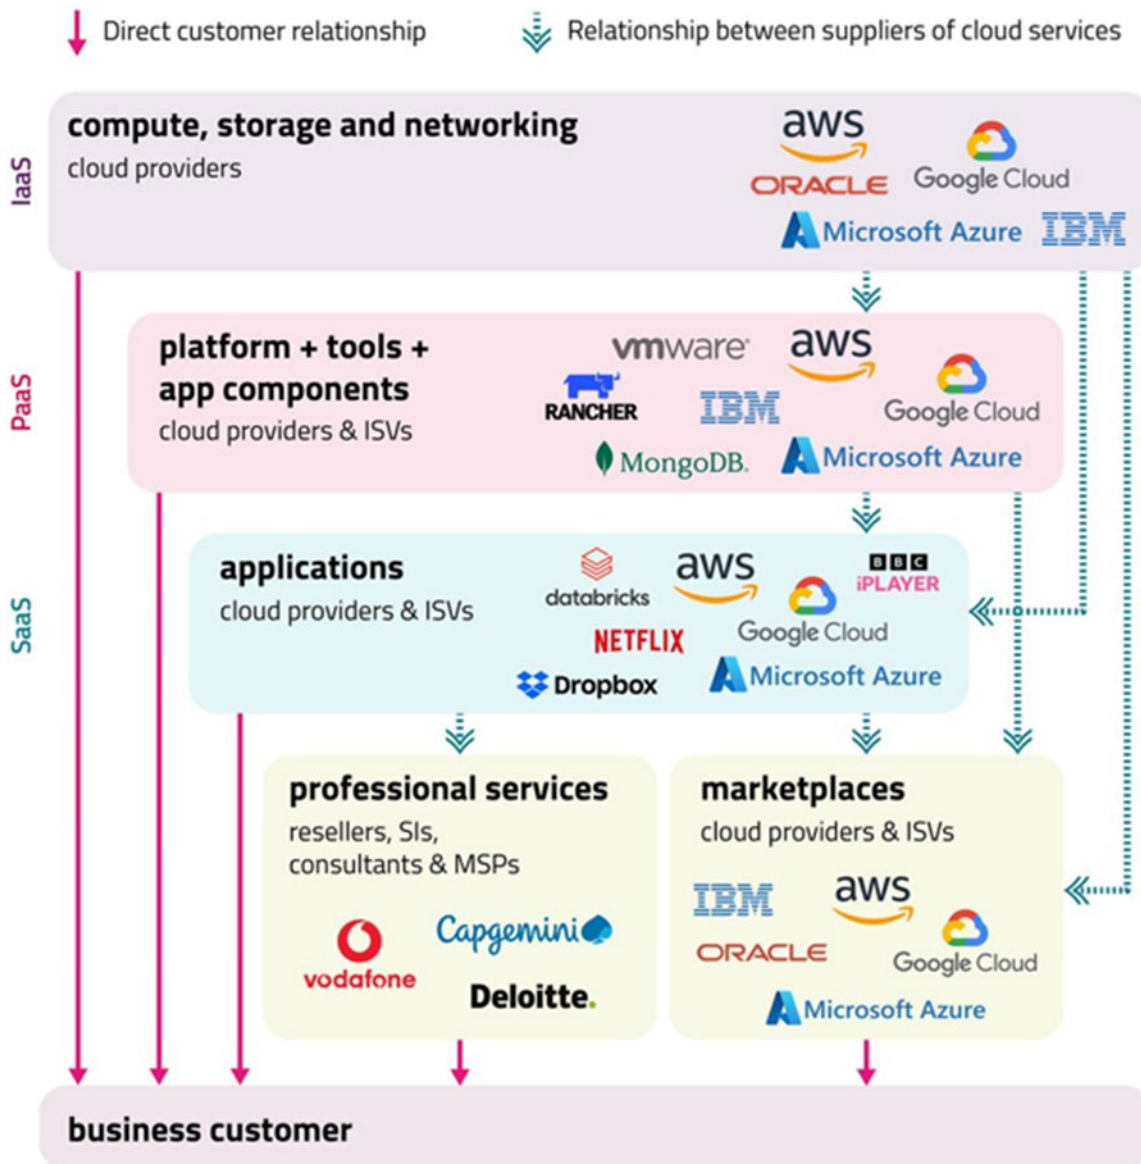
Market players

- 3.43 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:
- cloud providers, who are usually present at all levels of the cloud stack (IaaS, PaaS, SaaS);⁵⁸
 - independent software vendors (ISVs), who usually do not own their own infrastructure and are present at only one or two levels of the cloud stack (PaaS and/or SaaS); and
 - suppliers of professional services, who provide customers access to cloud services and/or support for using cloud services.
- 3.44 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.

⁵⁷ Flexera, 2023. 2023 State of the Cloud Report, page 18.

⁵⁸ 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

3.45 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

- 3.46 In the UK there are three hyperscalers - Amazon, Microsoft and Google.⁵⁹ 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.
- 3.47 These three hyperscalers are the main suppliers of public cloud infrastructure services in the UK. They collectively account for [3<] [70% to 80%] of total revenues generated from the supply of public cloud infrastructure services (IaaS and PaaS).⁶⁰ They also offer a wide portfolio of SaaS services.

Amazon

- 3.48 Amazon operates through its subsidiary AWS, which is considered as the overall market leader in cloud. We estimate that, as of 2021, AWS accounted for around [3<] [30% to 40%] of the UK's public cloud infrastructure revenues.
- 3.49 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.⁶¹ AWS currently offers over 210 cloud infrastructure services, that are organised into several product categories (e.g. compute, analytics, etc.).⁶²
- 3.50 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.⁶³ Amazon uses AWS to provide solutions across its businesses, for example, Prime Video uses AWS for compute, database and other services.⁶⁴ Amazon categorises its overall operations into three segments: AWS, North America and International.⁶⁵ While AWS is Amazon's smallest operating segment by revenue, representing about 16% of Amazon's revenue in 2022, it was Amazon's only profitable segment.⁶⁶

Microsoft

- 3.51 Microsoft is the second largest cloud provider and a close competitor to AWS. We estimate that, as of 2021, Microsoft accounted for around [3<] [30% to 40%] of the UK's public cloud infrastructure revenues.

⁵⁹ 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.

⁶⁰ 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.

⁶¹ AWS website. [About AWS](#) [accessed 19 January 2023].

⁶² Ofcom analysis of IaaS and PaaS products listed on AWS's website.

⁶³ AWS website. [What We Do](#) [accessed 19 January 2023].

⁶⁴ AWS website. [Prime Video Boosts Scale and Resilience](#) [accessed 19 January 2023].

⁶⁵ The 'North America' and 'International' segments largely consist of revenues from retail sales of consumer products.

⁶⁶ Ofcom analysis based on AWS's published financial statements. [Amazon.com announces fourth quarter results](#) [accessed 27 February 2023].

- 3.52 Before entering into 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 into the cloud computing market in 2008, with a PaaS offering which enabled developers to deploy applications in the cloud.⁶⁷ 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.⁶⁸ Microsoft currently offers over 200 cloud infrastructure services, across several product categories (e.g. compute, IoT).⁶⁹
- 3.53 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).⁷⁰ Microsoft's Intelligent Cloud segment includes Azure which provides cloud infrastructure and platform services (IaaS and PaaS). In the year to June 2022, Intelligent Cloud was Microsoft's largest operating segment, representing 38% of revenues and 39% of profit.⁷¹
- 3.54 Microsoft also uses Azure to provide a range of cloud services, such as parts of Office 365 and Bing Search, across its other operating segments.⁷² 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 2022, Microsoft Cloud accounted for 46% of Microsoft revenue.⁷³ 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

- 3.55 Google is the third largest cloud provider. 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 2021, Google accounted for approximately [§<] [5% to 10%] of the UK's public cloud infrastructure revenues.
- 3.56 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).⁷⁴ This platform was initially only made available to developers and was later made available as an

⁶⁷ Microsoft website. [About Microsoft](#) [accessed 6 March 2023]; and [Microsoft launches Windows Azure](#) [accessed 2 February 2023].

⁶⁸ Microsoft website. [Windows Azure Platform Now Generally Available](#) [accessed 6 March 2023]; [Office expands to the cloud](#) [accessed 2 February 2022]; and [The History of Microsoft Azure](#) [accessed 2 February 2022].

⁶⁹ Ofcom analysis of IaaS and PaaS products listed on Microsoft Azure's website.

⁷⁰ Microsoft website. [Segment Information](#) [accessed 2 February 2022].

⁷¹ Ofcom analysis based on Microsoft's published financial statements. [Microsoft 2022 10-K](#) [accessed 27 February 2022].

⁷² ZDNET, 2021. [Microsoft moves closer to running all of its own services on Azure](#) [accessed 2 February 2022].

⁷³ Ofcom analysis based on Microsoft's financial statements. [Microsoft 2022 10-K](#) [accessed 27 February 2022].

⁷⁴ Google website. [Google Cloud Platform](#) [accessed 6 March 2023].

official fully supported product in 2011.⁷⁵ 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).⁷⁶

- 3.57 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 Bets⁷⁷ (combination of all other services).⁷⁸ Google Cloud has so far been loss-making, and represented around 10% of Alphabet's total revenue in 2022, despite it being the fastest-growing segment of Alphabet in that year.⁷⁹
- 3.58 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.⁸⁰ 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.

Smaller cloud providers

- 3.59 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 that we collectively refer to as 'smaller cloud providers'. In 2021, the smaller cloud providers represented around [20% to 30%] of revenues associated with the supply of public cloud infrastructure services in the UK.⁸¹ These providers compete with the hyperscalers by offering services across the cloud stack, as well as partner with the hyperscalers by offering complementary services. We illustrate further the role of smaller cloud providers, using IBM and Oracle as examples, below.

⁷⁵ Cloud Guru, 2018. [The History of Google Cloud Platform](#) [accessed 6 March 2023].

⁷⁶ Ofcom analysis of IaaS and PaaS products listed on Google Cloud's website.

⁷⁷ According to Alphabet, revenues from Other Bets are generated primarily from the sale of health technology and internet services.

⁷⁸ Alphabet website. [Alphabet Announces First Quarter 2022 Results](#) [accessed 2 February 2023].

⁷⁹ Ofcom analysis based on Alphabet's published financial statements. [2022 Annual report](#) [accessed 27 February 2023].

⁸⁰ Google's website. [What is Cloud Computing?](#) [accessed 2 February 2022]. DCD, 2021. [Google to migrate parts of YouTube to Google Cloud](#) [accessed 2 February 2023].

⁸¹ 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.

IBM

IBM first started providing cloud services around 2008, with the launch of a SaaS offering. It later expanded to provide IaaS and PaaS.⁸² 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.⁸³

IBM currently offers a range of services (e.g. compute, AI and machine learning).⁸⁴ 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.⁸⁵ IBM's customers tend to be mainly large and mid-size enterprises.⁸⁶ 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.⁸⁷

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.⁸⁸

Oracle appear to have a focus on enterprise customers, hybrid and multi-cloud offerings.⁸⁹ Oracle has agreements with Microsoft, which allow its services, including its database offerings, to run on Microsoft's platform with "seamless interoperability".⁹⁰

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).⁹¹ Oracle identify its cloud to be "the first and only sovereign, dedicated dual-region cloud for UK Government and Defence customers".⁹²

Independent software vendors (ISVs)

3.60 ISVs are suppliers of cloud services, typically PaaS and/or SaaS, that do not usually own the underlying infrastructure.⁹³ ISVs and cloud providers interact in a number of different ways:

⁸² Datamation, 2017. [What is IBM Cloud?](#) [accessed 3 February 2023].

⁸³ Tech Monitor, 2022. [What is IBM known for?](#) [accessed 3 February 2023].

⁸⁴ IBM website. [IBM Cloud products](#) [accessed 3 February 2023].

⁸⁵ RedHat, 2019. Press release. [IBM Closes Landmark Acquisition of Red Hat](#) [accessed 3 February 2023].

⁸⁶ [X].

⁸⁷ IBM website. [Regulated workloads](#) [accessed 3 February 2023].

⁸⁸ Oracle website. [About Oracle](#) and [Oracle Cloud Infrastructure Platform](#) [accessed 3 February 2023].

⁸⁹ Oracle website. [Oracle's distinct approach on hybrid and multi-cloud](#) [accessed 30 March 2023].

⁹⁰ Oracle website. [Oracle Interconnect for Azure](#) [accessed 3 February 2023].

- a) ISVs rely on cloud providers as 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);⁹⁴
- 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 vertically integrated 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 ISV 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.

3.61 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.⁹⁵ VMware's strategy is focused on enabling multi-cloud environments.⁹⁶

VMware partner 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).⁹⁷ 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 appear to compete with cloud providers' own cloud native virtualisation solutions.⁹⁸

⁹¹ Techzine, 11 July 2022. [Oracle expands OCI with sovereign cloud regions inside EU](#) [accessed 30 March 2023].

⁹² Oracle website. [Oracle Cloud for UK Government & Defence](#) [accessed 3 February 2023].

⁹³ Although, we note that there may be exceptions to this. For example, Salesforce is an ISV that also operates its own infrastructure.

⁹⁴ We recognise that the PaaS/SaaS categorisation can be arbitrary.

⁹⁵ VMware website. [VMware Timeline](#) [accessed 3 February 2023].

⁹⁶ VMware website. [Multi-Cloud Environments](#) [accessed 3 February 2023].

⁹⁷ AWS website. [VMware Cloud on AWS](#) [accessed 3 February 2023].

⁹⁸ VMware response dated 9 December 2022 to the s.174 notice dated 27 October 2022, Part A, question 14, para 14.2.

MongoDB

MongoDB was founded in 2007. It specialises in database management and document databases.⁹⁹ 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.¹⁰⁰

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. [§<].¹⁰¹ 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.¹⁰²

Snowflake

Snowflake was founded in 2012.¹⁰³ 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.¹⁰⁴

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.¹⁰⁵ Snowflake also identify the hyperscalers as its competitors.¹⁰⁶ 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).

⁹⁹ MongoDB website. [Our Mission](#) [accessed 3 February 2023].

¹⁰⁰ MongoDB website. [MongoDB Atlas](#) [accessed 6 March 2023]; and [Advantages of MongoDB](#) [accessed 3 February 2023].

¹⁰¹ [§<].

¹⁰² Redburn, 2022. Cloud computing report, page 94.

¹⁰³ Techstory, 2020. [The Snowflake story](#) [accessed 23 March 2023].

¹⁰⁴ Snowflake website. [Enabling the data cloud with Snowflake](#) [accessed 3 February 2023].

¹⁰⁵ Snowflake website. [Cloud partners](#) [accessed 23 March 2023].

¹⁰⁶ Snowflake response dated 7 December 2022 to the s.174 notice dated 26 October 2022, question 12, page 11.

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.¹⁰⁷ YugabyteDB is available to customers as both fully managed and self-managed services.¹⁰⁸

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 with hyperscalers to win clients, and also compete with the database offerings of hyperscalers.¹⁰⁹

Suppliers of professional services

3.62 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.¹¹⁰ 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.

3.63 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

¹⁰⁷ Yugabyte website. [About Yugabyte](#) [accessed 28 February].

¹⁰⁸ Yugabyte website. [YugabyteDB Deployment Options](#) [accessed 28 February].

¹⁰⁹ Ofcom meeting with Yugabyte on 22 November 2022.

¹¹⁰ 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.

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).

- 3.64 We observe that most suppliers (e.g. Accenture,¹¹¹ Capgemini¹¹²) 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.¹¹³ We are also aware of cases where some suppliers indicate that they compete with the hyperscalers. For example, [redacted].¹¹⁴ However, we note such cases appear overall to be limited.

Cloud services in the telecoms and broadcasting sectors

Telecoms

- 3.65 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.
- 3.66 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

- 3.67 UK telecoms providers told us that some of their network functions¹¹⁵ have started moving onto the cloud. Where telecoms providers are using the cloud to host these network functions, they are broadly using private cloud.¹¹⁶ BT told us it thinks telecoms providers who have invested in their private clouds are unlikely to migrate fully to the public cloud,¹¹⁷ suggesting hybrid cloud architecture will be used in the future. This view was supported by [redacted] [redacted] and [redacted].¹¹⁸ [redacted] also told us that they have plans to begin testing of some 'non-critical' network functions in the public cloud.¹¹⁹
- 3.68 UK telecoms providers are not using public cloud for their network functions, and instead have maintained substantial private clouds. Many other cloud users use the public cloud for core workloads specific to their sectors.¹²⁰ Telecoms providers highlighted some

¹¹¹ Accenture website. [Cloud services](#) [accessed 22 March 2023].

¹¹² Capgemini website. [Technology partners](#) [accessed 22 March 2023].

¹¹³ Ofcom meeting with Capgemini on 1 November 2022.

¹¹⁴ Ofcom meeting with [redacted].

¹¹⁵ Network function is a component of telecom networks that delivers a specific function (e.g. router, switch, load balancer, firewall).

¹¹⁶ Capgemini, 2023. [Networks on Cloud: a clear advantage](#), page 6 [accessed 8 March 2023].

¹¹⁷ BT response to the CFI, page 6, paragraph 2.

¹¹⁸ Ofcom / [redacted] meeting, [redacted], Ofcom / [redacted] meeting, [redacted] and Ofcom / [redacted] meeting, [redacted].

¹¹⁹ Ofcom / [redacted] meeting, [redacted] and [redacted] response to Ofcom disclosure document, [redacted].

¹²⁰ For example: Monzo, 2021. [An introduction to Monzo's data stack](#). [accessed 17 March 2023].

drawbacks to using public cloud for telecoms specific workloads including a lack of control, a risk to resiliency and uncertain costs.¹²¹ For example, in their CFI response, BT 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.¹²²

- 3.69 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 [redacted] 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.¹²³
- 3.70 We have seen recent examples of new cloud products being developed and launched specifically for telecoms providers.¹²⁴ As the suite of products targeted at telecoms providers grows, we may see increased use of public cloud services for some network functions.¹²⁵
- 3.71 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.¹²⁶

Multi-access edge computing (MEC)

- 3.72 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.¹²⁷ BT 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.¹²⁸ 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.¹²⁹
- 3.73 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

¹²¹ Ofcom / [redacted] meeting, [redacted] and Ofcom / [redacted] meeting, [redacted].

¹²² BT response to the CFI, page 6, paragraph 1.

¹²³ Ofcom / [redacted] meeting, [redacted].

¹²⁴ Examples include [AWS Telco Network Builder](#), [Snowflake Telecom Data Cloud](#) and [Microsoft telecom cloud tools](#) [accessed 28 February 2023].

¹²⁵ 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 7 March 2023].

¹²⁶ ABI research, 2023 <https://go.abiresearch.com/lp-public-private-or-hybrid-cloud-for-5g-core-network-deployments> [accessed 8 March 2023].

¹²⁷ Vodafone website. [Multi-Access Edge Computing](#) [accessed 2 March 2023].

¹²⁸ BT Wholesale website. [BT uses Amazon Web Services \(AWS\) Wavelength](#) [accessed 21 March 2023].

¹²⁹ Ofcom / [redacted] meeting, [redacted] and Ofcom / [redacted] meeting, [redacted].

bandwidth.¹³⁰ 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

3.74 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.¹³¹ 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 (siloes multi-cloud). The drivers for moving to 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

3.75 Telecoms providers sometimes have a multi-layered relationship with the hyperscalers. In addition to being customers they are, in some instances, partners of hyperscalers.¹³² 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).

3.76 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.¹³³

Broadcasting

3.77 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).

3.78 We received evidence that some broadcasters use a combination of public and private cloud offerings. [redacted] told us that key drivers for using private cloud include low latency and

¹³⁰ Heavy Reading, 2022. [5G Network Strategies: Operator Survey 2022](#) [accessed 2 March 2023].

¹³¹ BT response to CFI, page 7, paragraph 2, [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted] and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹³² 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](#) [accessed 16 March 2023].

¹³³ AWS website. [AWS Private 5G](#) [accessed 16 February 2023].

information security policy requirements.¹³⁴ However, [redacted] noted that they do not use private cloud, as they find public cloud offerings to be sufficient in areas such as security and resilience.¹³⁵

- 3.79 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.
- 3.80 [redacted] 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 argued 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.¹³⁶
- 3.81 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.

¹³⁴ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹³⁵ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹³⁶ [redacted] response dated [redacted] to the CFI, page [redacted], paragraph [redacted].

4. How competition works

Introduction

4.1 This section looks at how competition works in the cloud 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 the extent to which customer preferences are reflected in the range and magnitude of benefits cloud providers typically offer to attract and acquire customers. 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. And third, we consider key market outcomes, including the positioning of the key providers in the market. We look at the market shares of UK cloud services, analyse the profitability of the key providers, and consider the service offerings and key capabilities of each hyperscaler.

How customers choose and buy cloud services

- 4.2 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.
- 4.3 We acknowledge that 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

- 4.4 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,¹³⁷ Oracle 2023 Report,¹³⁸ and others also often classify cloud customers by industry and size, as well as geography.
- 4.5 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. We generally observe that cloud providers tend to have a number of large, high-spend customers (e.g.

¹³⁷ Flexera, 2023. 2023 State of the Cloud Report.

¹³⁸ 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 23 March 2023].

enterprises or government departments) that account for a small proportion of their customer base but a large proportion of their revenues. 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.¹³⁹

- 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.
- 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”.¹⁴⁰ 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.

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

How customers buy cloud services

Purchase channels

4.7 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 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%).

¹³⁹ [redacted] response dated [redacted] to our follow-up email concerning the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to our follow-up email concerning the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁴⁰ Context Consulting research report, Slide 20.

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

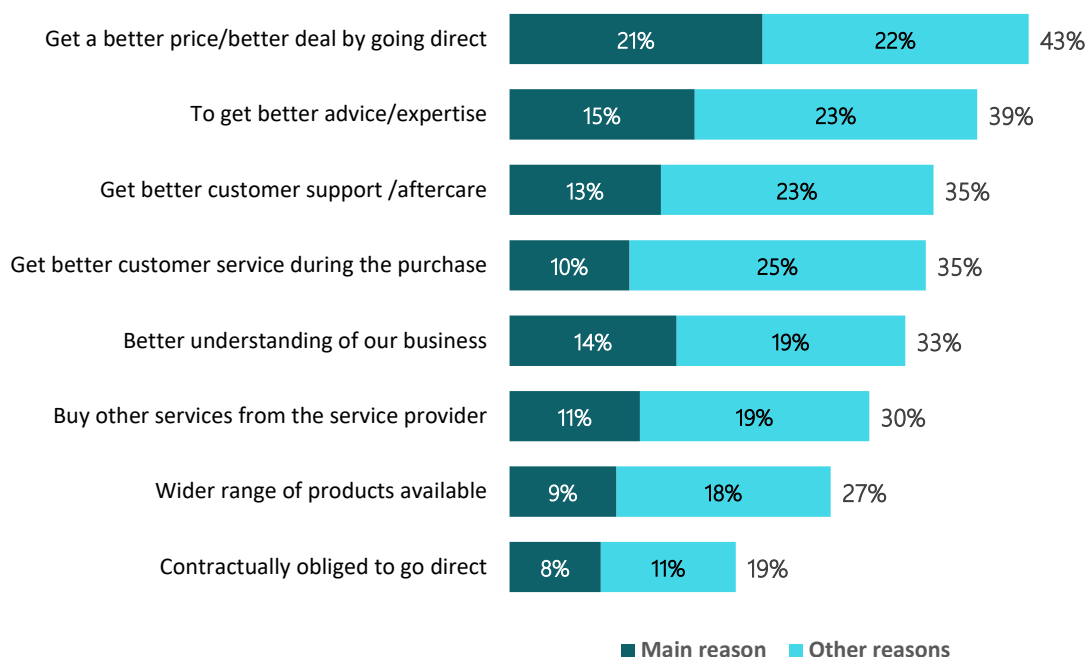
| Channels of purchase ¹⁴¹ | IaaS users | PaaS users |
|-------------------------------------|------------|------------|
| Direct from hyperscaler | 51% | 38% |
| Via a telecoms provider | 33% | 36% |
| Via a service integrator | 32% | 30% |
| Via a managed service provider | 32% | 36% |
| Via other provider | 0% | 1% |

Source: Context Consulting research data tables, Q33.

- 4.8 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%).
- 4.9 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.

¹⁴¹ These responses are not mutually exclusive.

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



Source: Context Consulting research report, Slide 87.

- 4.10 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%).¹⁴² 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.
- 4.11 Many of the large customers that responded to our statutory information requests and customer questionnaire told us they are not using intermediaries ([§]).¹⁴³ 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.
- 4.12 [§] 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

¹⁴² Context Consulting research report, Slide 87.

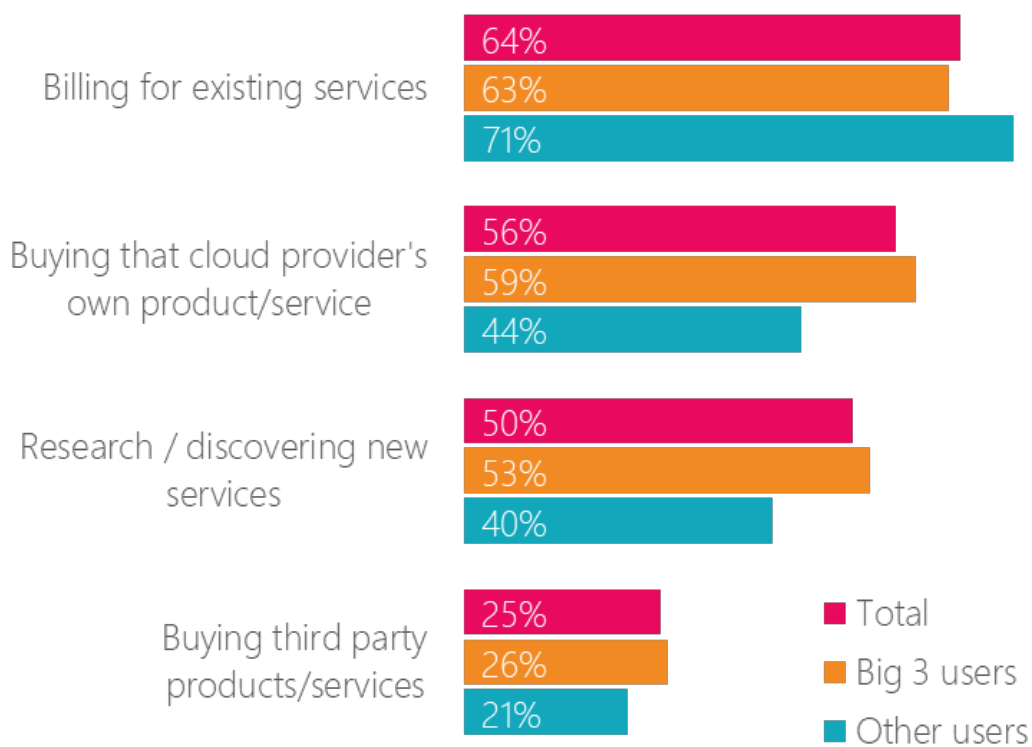
¹⁴³ [§] response dated [§] to the s.174 notice dated [§], question [§]; [§] response dated [§] to the s.174 notice dated [§], question [§]; [§] response dated [§] to the s.174 notice dated [§], question [§]; [§] response dated [§] to the s.174 notice dated [§], question [§]. [§] response dated [§] to 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 [§].

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.¹⁴⁴

Marketplace use

4.13 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%).¹⁴⁵

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



Source: Context Consulting research data tables, Q46.

4.14 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%).

4.15 Buying third-party products and services was mentioned by 25% of those who use marketplaces (or 13% of all IaaS/PaaS users).¹⁴⁶ Buying third-party products and services is relatively more important for larger companies with 2500+ employees (at 31% of those

¹⁴⁴ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁴⁵ Context Consulting research data tables, Q45.

¹⁴⁶ Analysis of Context Consulting research data, Q45 & 46.

who use marketplaces), companies older than 20 years (31%), those in public sector (40%), and AWS users (33%).

- 4.16 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 ([§]).¹⁴⁷ [§] told us it very rarely uses marketplaces because its internal processes for authorising such purchases are burdensome.¹⁴⁸ [§] 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.¹⁴⁹
- 4.17 At the same time, some other customers told us they use marketplaces ([§]),¹⁵⁰ mainly because spend through marketplaces is counted towards securing discounts, and billing and terms are simplified.

Contracts, pricing, negotiation

- 4.18 The vast majority of respondents in our market research had contracts with cloud providers, only about 17% of user-provider relationships¹⁵¹ were on a pay-as-you-go basis (i.e. respondents that do not have an ongoing contractual relationship). Out of all the cases when the contract length was provided (67% of responses), the most common contract lengths were 2 or 3 years, while 5-year contracts were common as well.
- 4.19 We asked the hyperscalers about the length of their spending commitments and agreements with customers. They told us that in 2022:
- [§]% of their total UK cloud 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;
 - [§]% of revenues were accounted for by customers with spending commitments and agreements lasting at least 5 years; and

¹⁴⁷ [§] 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 [§].

¹⁴⁸ [§] response dated [§] to the s.174 notice dated [§], question [§].

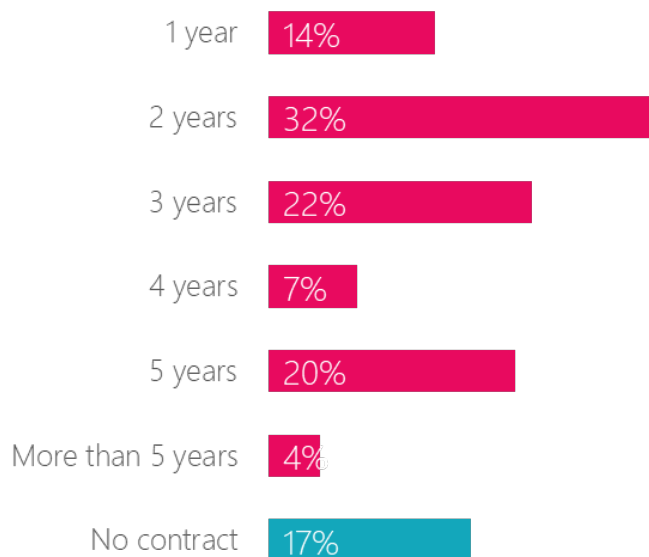
¹⁴⁹ [§] response dated [§] to the s.174 notice dated [§], question [§].

¹⁵⁰ [§] 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 [§].

¹⁵¹ 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.

- d) there was more variation in use of agreements longer than 5 years, which accounted for [X]% of revenues to [X]% of revenues.¹⁵²

Figure 4.4. Frequency of various lengths of contract, as reported in our market research



Source: Context Consulting research report, Slide 94.

Note: the percentages of different contract lengths are calculated on per-contract basis, not on per-respondent basis as one respondent can have several interactions with different providers. Responses “no contract” or “don’t know” were excluded from the denominator of the calculation. “No contract” percentage was calculated as number of “no contract” responses (283) divided by total sum of user-provider relationships (1,682). No contract refers to those that do not have an ongoing contractual relationship

- 4.20 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 (non-cloud) products** or services. Overall, about 4 in 10 “user-provider relationships” were reported to have other services bundled in, and in 56% of cases it was a cloud-only purchase.¹⁵³ Purchases from AWS were more likely to be stand-alone (62%) than from Microsoft (53%) or Google (52%).
- 4.21 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%).¹⁵⁴ 5% of respondents said they had to buy cloud services together with other services.
- 4.22 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 6% of cases there was a

¹⁵² One hyperscaler was only able to provide this information for 2021 rather than 2022. [X] response dated [X] to follow on questions dated [X], question [X]; [X] response dated [X] to follow on questions dated [X], question [X]; and [X] response dated [X] to follow on questions dated [X], question [X].

¹⁵³ Context Consulting research report, Slide 89.

¹⁵⁴ Context Consulting research report, Slide 89.

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.

- 4.23 In about half (52%) of cases the respondents found it 'quite easy' or 'very easy' to accurately **estimate future costs**. In 18% of cases, it was 'neither easy nor challenging', and in about 30% of cases it was somewhat or very challenging. The results for larger providers are similar to the average, with 55% of Microsoft customers saying it was very or quite easy, 48% for AWS and 57% for Google.
- 4.24 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. [X] 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.¹⁵⁵ [X] told us that there are limitations on the extent to which businesses of [X] 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, [X] suggested these are limited and based on spend levels. And while [X] 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.¹⁵⁶
- 4.25 [X] 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. [X] 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.¹⁵⁷
- 4.26 However, a greater number of customers [X] say they were able to negotiate bespoke agreements with cloud providers.¹⁵⁸ 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

- 4.27 Several cloud providers in response to our statutory information requests told us that tenders are often held by clients in the public sector where it can be a regulatory

¹⁵⁵ [X] response dated [X] to the s.174 notice dated [X], question [X].

¹⁵⁶ [X] response dated [X] to the s.174 notice dated [X], question [X].

¹⁵⁷ [X] response dated [X] to the s.174 notice dated [X], question [X].

¹⁵⁸ [X] response dated [X] to the s.174 notice dated [X], question [X]; [X] response dated [X] to the s.174 notice dated [X], question [X]; [X] response dated [X] to the s.174 notice dated [X], question [X]; [X] response dated [X] to the s.174 notice dated [X], question [X]; [X] response dated [X] to the s.174 notice dated [X], question [X]; [X] response dated [X] to the s.174 notice dated [X], question [X]; [X] response dated [X] to the s.174 notice dated [X], question [X].

requirement [redacted].¹⁵⁹ [redacted] also suggested it is more common for larger organisations to run competitive tenders.¹⁶⁰

- 4.28 Some customers in response to our statutory information requests told us they used tenders before ([redacted]), usually for limited specific purposes, while others ([redacted]) never used them.¹⁶¹

Contract renegotiation

- 4.29 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). We asked those who renegotiated their contract several questions about their renegotiation experience including whether they experienced a price rise.
- 4.30 Overall, in 47% of instances of renegotiation the respondents found the process easy. In 23% of renewals/renegotiations customers discovered that at least some of their preferred products were not available anymore, in 15% of cases the new contract was more restrictive, and in 6% of cases they were encouraged to buy more products than they needed.
- 4.31 In terms of individual respondents (and not user-provider relationships), 59% renewed or renegotiated at least one of their contracts, and out of those who renewed/renegotiated 56% experienced a price rise.¹⁶²
- 4.32 While the mean average price rise was 20% across all 425 cases of price rises in the market research, the median price rise was 10%, and in 10% of cases the price rise was 50% or more.¹⁶³ At the same time, on the lower end of price increases, 10% of customers saw price rises of 4% or lower. The following chart shows the distribution of price rises excluding cases when the rise was above 100% (to make the chart more readable).

¹⁵⁹[redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted] question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

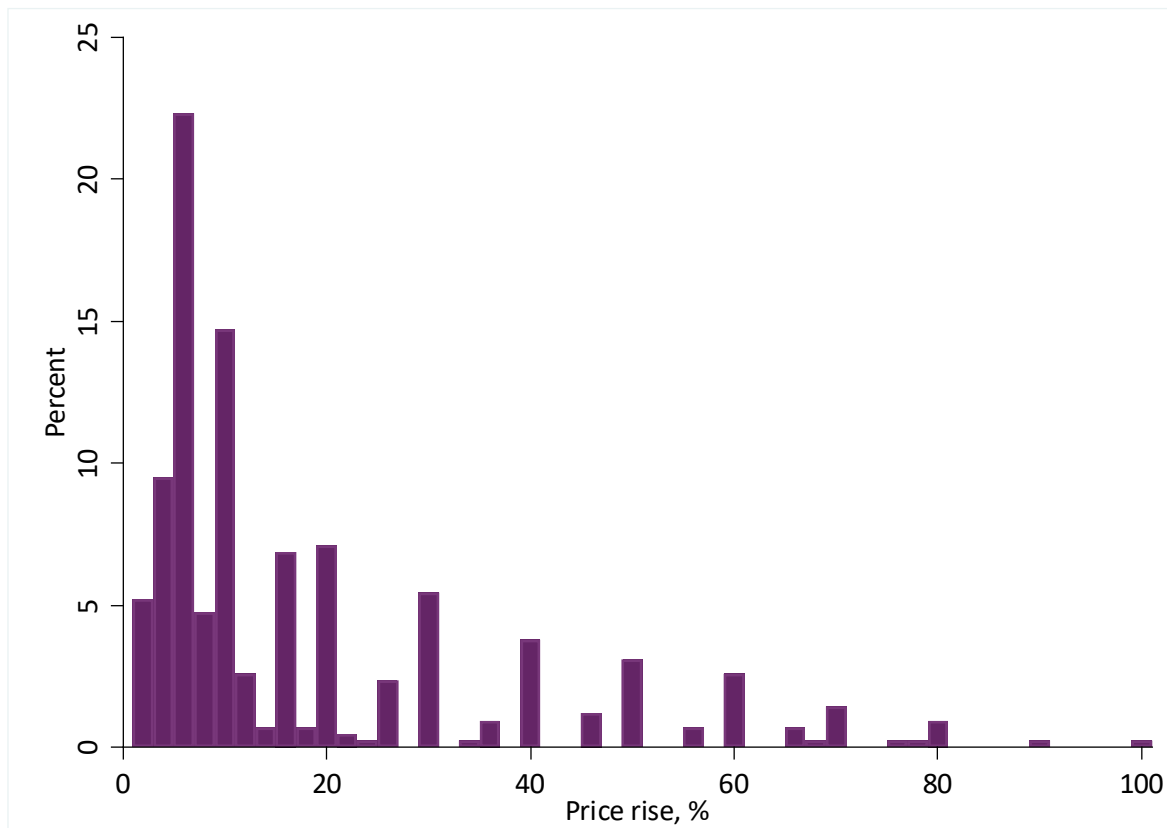
¹⁶⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁶¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁶² Analysis of Context Consulting research data, Q42 & 43.

¹⁶³ Price rises above 100% were reported in 4 cases out of 425.

Figure 4.5. Distribution of price rises experienced, as reported in our market research (excluding rises of more than 100%)



Source: Analysis of Context Consulting research data, Q44.

- 4.33 For the hyperscalers the picture was close to the overall numbers across all providers. For Microsoft the average price rise was 19% and median was 10%, for AWS – 20% and 10% respectively, and for Google – 24% and 15% respectively. The lower 10th percentiles were 3%, 4% and 4% (for Microsoft, AWS and Google, respectively), and the upper 10th percentiles were 40%, 55% and 55%.
- 4.34 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. It is common to see customers report a higher discount in exchange for agreeing to a higher committed spend. Renegotiation is not always precipitated by the approaching end of contract, but can be prompted by changes in customer’s needs, or the customer seeking better terms in light of its increased use of the cloud.

How customers choose between different providers

- 4.35 The key decision point when choosing a cloud provider is when the customer first migrates into the cloud, but for some customers there can be further decisions 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.

4.36 Based on the information we have gathered in the first half of this study, we have identified the following factors that can determine customers' choice of cloud provider. 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.¹⁶⁴ 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.¹⁶⁵ 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.¹⁶⁶
- 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

¹⁶⁴ Context Consulting research report, Slide 68.

¹⁶⁵ Flexera, 2023. Flexera 2023 State of the Cloud report. Flexera suggest that respondents anticipate organisational spend on public cloud to grow by 30% in the next year, p. 41.

¹⁶⁶ "Offered the best price" comes only 8th on the list of factors for provider choice in our market research, mentioned in 24% of cases.

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 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) **Other factors.** 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. Furthermore, considerations such as 'increasing focus on environmental impact' may lead customers to give weight to the sustainability credentials of different cloud providers.

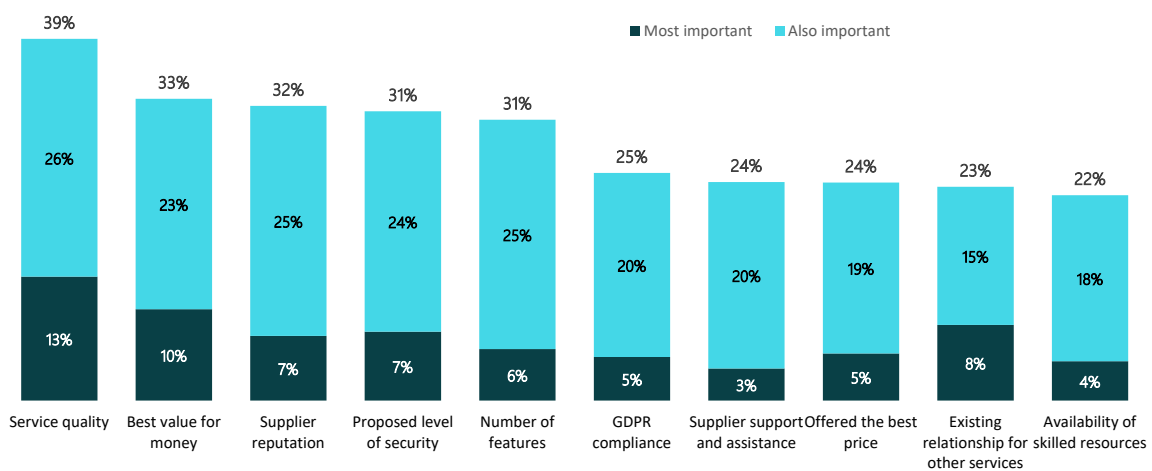
Some customers use a multi-stage decision process

- 4.37 While there are several factors customers consider when choosing a provider, 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.
- 4.38 Other customers did not describe choosing a provider as a multi-stage process and consider all relevant factors in the round when making a decision.

Relative importance of factors

4.39 Figure 4.6 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. Across all cases, the top 5 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 reason, with “existing relationship for other services” moving up to the third place of most important reasons.

Figure 4.6. Reasons for choosing a provider, as reported in our market research



Source: Context Consulting research report, Slide 68.

4.40 Respondents were asked about the reasons for choosing a provider for each provider they were using, Figure 4.7 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.

4.41 ‘Existing relationship for other services’ is relatively more important for Microsoft Azure and BT customers, while is not in the top 6 reasons for other major providers. Supplier support are among the top factors for Google, IBM and BT. Only for Oracle customers did ‘availability of skilled resources’ reach the top 6 reasons.

Figure 4.7. 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.

4.42 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.¹⁶⁷ 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

4.43 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 from a cloud provider.

4.44 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. Below, we examine the set of customer needs that are likely to be met by each of these deployment models and the extent to which customers may consider these deployment models as substitutable.

¹⁶⁷ Context Consulting research report, Slide 132.

- 4.45 We recognise that there may be some substitutability between public cloud and the other deployment models. In particular, some cloud providers (e.g. [redacted]) indicate that they do not distinguish their products across the different cloud deployments that they cater to (public, (virtual) private and hybrid) and that customers have the ability to deploy relevant products in any of these models.¹⁶⁸ Also for certain use cases, such as storing data, some cloud providers (e.g. [redacted]) suggest that customers can use cloud and on-premises IT solutions to run the same workloads and address the same requirements.¹⁶⁹
- 4.46 We think that 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 customers. We observe that customers:
- a) use public cloud to often meet requirements such as scalability and/or innovation (e.g. Sainsbury, ITV, Netflix).¹⁷⁰ These advantages are a key reason for why customers are increasingly migrating workloads from on-premises IT to public cloud. Both Sainsbury and ITV suggest that the use of public cloud gives them the flexibility to respond to customer needs and scale up in times of increased demand.¹⁷¹ In particular, Sainsbury note that they wouldn't 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.¹⁷² 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 "costs per streaming start" has reduced to a fraction of the costs previously incurred when relying on their in-house data centres;¹⁷³
 - b) use private cloud to meet specific requirements around latency, security, resilience and/or regulatory compliance (e.g. [redacted]).¹⁷⁴ 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;

¹⁶⁸ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁶⁹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁷⁰ Sainsbury 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 13 February 2023].

¹⁷¹ Sainsbury 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.

¹⁷² Sainsbury response dated 8 December 2022 to our customer questionnaire, question 19.

¹⁷³ Netflix, 2016. [Completing the Netflix Cloud Migration](#) [accessed 13 February 2023].

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

- 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 cloud (e.g. [redacted]¹⁷⁵); and
 - d) use on-premises solutions where there is a need to maintain control over physical hardware for security and regulatory requirements (e.g. [redacted]¹⁷⁶). For example, 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.¹⁷⁷
- 4.47 We observe that the ability for customers to switch between these deployment models will ultimately depend on the extent to which a given application is designed to take advantage of the unique environment provided by that deployment model. For example, as part of their current strategy, some customers (e.g. [redacted]) suggest that they expect that some of their workloads will always remain on private cloud or on-premises due to performance, security and regulatory requirements.¹⁷⁸ At the same time, we are aware of other customers who have migrated workloads from private to public cloud, and consider themselves unlikely to return to private cloud.
- 4.48 Overall, we consider that each deployment model meets a distinct set of customer needs. We also consider that once a customer has decided to migrate a workload into the public cloud, alternative deployment models are likely to be a weak alternative for these workloads. Similarly, our engagement with stakeholders suggests that some workloads are not well-suited for use in the public cloud, and the latter is likely to pose a weak alternative for these workloads currently run in a private cloud. We therefore consider it appropriate to focus our competitive assessment in remaining sections on providers of public cloud infrastructure services (unless otherwise stated), without further considering the potential competitive constraints imposed by other deployment models.

How providers compete

- 4.49 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).
- 4.50 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

¹⁷⁵ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁷⁶ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁷⁷ TechCrunch, 2019. [DropBox infrastructure continues to evolve](#) [accessed 13 February 2023].

¹⁷⁸ [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

providers and ISVs use, before discussing the role of ecosystems in how providers compete in cloud.

Pricing and costs

4.51 As identified earlier, 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.¹⁷⁹

- a) **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 certain monetary value of services over a particular contract length) can either apply to individual products or families of products (these tend to be generally available to all customers), or to a customer's total spend with a cloud provider (these tend to be privately negotiated). [redacted] suggest 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.¹⁸⁰ 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, pay-as-you-go prices.¹⁸¹
- b) **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.¹⁸² [redacted] and [redacted] suggest they provide credits so that new customers can test and explore their cloud services.¹⁸³ 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'.¹⁸⁴ We are also aware that cloud credits may be offered to

¹⁷⁹ 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.

¹⁸⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁸¹ AWS website. [Amazon EC2 Spot Instances](#) [accessed 7 March 2023]; and Microsoft website. [Azure Spot Virtual Machines](#) [accessed 7 March 2023].

¹⁸² Microsoft website. [Azure free account](#) [accessed 17 February 2023]; and IBM website. [IBM Cloud free tier](#) [accessed 17 February 2023].

¹⁸³ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁸⁴ AWS website. [AWS activate](#) [accessed 17 February 2023]; Microsoft website. [Unlocking Azure credits as your start-up grows](#) [accessed 17 February 2023]; Google website. [Google for start-ups cloud program](#) [accessed 17 February 2023].

customers on a case-by-case basis and therefore there may be scope for customers to negotiate and benefit from higher credit amounts.¹⁸⁵

- c) **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.¹⁸⁶ 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.¹⁸⁷

Quality and range of services

4.52 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.

- a) **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.¹⁸⁸
- b) **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”.¹⁸⁹ By continuing to innovate, cloud providers can improve their performance and technical capabilities, therefore enabling them to provide a better quality of service for customers.
- c) **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 as this can determine their ability to switch and or build their preferred cloud architecture by combining services from multiple clouds.

Geographic coverage

4.53 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

¹⁸⁵ Such credit offerings may also be linked to spend commitments. [redacted] response dated [redacted] to the s.174 notice [redacted], question [redacted].

¹⁸⁶ AWS website. [AWS Free Tier](#) [accessed 17 February 2023]. We note that cloud providers may also additionally provide certain cloud services as free of charge always (permanent free tier).

¹⁸⁷ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁸⁸ Ofcom analysis of [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

¹⁸⁹ Ofcom / Capgemini meeting, 1 November 2022.

to other factors.¹⁹⁰ 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.¹⁹¹
- 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 data centres 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.¹⁹² Companies operating in different regions across the world may also look for a provider that has data centres 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).

Sales channels

4.54 Table 4.8 sets out the proportion of sales by sales route for some cloud providers and ISVs. 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 Microsoft and VMware). Sales through first- and third-party marketplaces are in general a small proportion of providers' total sales revenue.

¹⁹⁰ "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.

¹⁹¹ 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.

¹⁹² 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.

Table 4.8: Estimated proportion of UK sales (by revenue) through different sales channels, 2021

| | Direct sales | Suppliers of professional services | Marketplaces |
|------------------|---------------|------------------------------------|--------------|
| AWS | [X]% | [X]% | [X]% |
| Google | [X]% | [X]% | [X]% |
| Microsoft | [X] [20-30%] | [X] [70-80%] | [X] [0-5%] |
| IBM | [X] [90-100%] | [X] [5-10%] | [X] [0-5%] |
| Oracle | [X] [80-90%] | [X] [10-20%] | [X] [0-5%] |
| OVHcloud | [X]% | [X]% | [X]% |
| MongoDB | [X] [70-80%] | [X] [0-5%] | [X] [20-30%] |
| Snowflake | [X] [80-90%] | [X] [0-5%] | [X] [10-20%] |
| VMware | [X] [5-10%] | [X] [90-100%] | [X] [0-5%] |

Sources: AWS response dated 9 December 2022 to s.174 notice of 24 October 2022, Part B Q5, and Annex 2; Google response dated 16 December 2022 to s.174 notice of 26 October 2022, Part B Q5, Annex 2; IBM response dated 6 December 2022 to s.174 notice of 25 October 2022, Part B Q5, Annex 2; Microsoft response dated 9 December 2022 to s.174 notice of 21 October 2022, Part B Q5 and Confidential Annex B5; MongoDB response dated 16 December 2022 to s.174 request of 27 October 2022, Part B Q5, Annex: Part B Cloud Services Template updated; Oracle response dated 16 December 2022 to s.174 notice of 31 October 2022, Part B Q5, Annex 2; OVHcloud response dated 12 December 2022 to s.174 notice of 27 October 2022, Part B Q5, Annex B; Snowflake response dated 7 December 2022 to s.174 request of 26 October 2022, Part B Q5 & Appendix 1; VMware response dated 14 December 2022 to s.174 notice of 27 October 2022, Part B Q5 and Annex. Notes: [X].

4.55 Overall, we estimate that less than 10% of total ISV sales (PaaS and SaaS) in the UK are transacted via a hyperscaler marketplace.¹⁹³

Ecosystems in cloud services

4.56 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 can be seen in how some cloud providers operate within the cloud market, making ecosystems a relevant framework through which to examine competition within

¹⁹³ 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.

the market. There are several cloud providers that offer ecosystems, but our focus here is on the hyperscalers as the main providers in the UK.

- 4.57 For example, 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.

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

- 4.58 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 [redacted] [70% to 80%] of UK IaaS revenues in 2021 and [redacted] [40% to 50%] of UK PaaS revenues.¹⁹⁴ 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.¹⁹⁵
- 4.59 Customers have told us they value the range of products offered by the hyperscalers,¹⁹⁶ and this is reflected in customer purchase data, with the average hyperscaler customer purchasing multiple services.¹⁹⁷ 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,¹⁹⁸ streamlined procurement processes,¹⁹⁹ and can simplify customers’ recruitment and training processes as their staff only need to specialise in a single cloud ecosystem.²⁰⁰
- 4.60 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

¹⁹⁴ 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.

¹⁹⁵ Ofcom analysis of data from IDC. This is set out in more detail in the below subsection on UK shares of supply.

¹⁹⁶ [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; and [redacted] response dated [redacted] to our customer questionnaire, question [redacted]. [redacted] response dated [redacted] to s.174 request of [redacted], [redacted]; [redacted] response dated [redacted] to s.174 request of [redacted], [redacted]; and [redacted] response dated [redacted] to s.174 request of [redacted], [redacted].

¹⁹⁷ [redacted] response dated [redacted] to s.174 request of [redacted], [redacted]; and [redacted] response dated [redacted] to s.174 request of [redacted], [redacted].

¹⁹⁸ [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; and [redacted] response dated [redacted] to our customer questionnaire, question [redacted].

¹⁹⁹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

²⁰⁰ [redacted] response dated [redacted] to the CFI, page [redacted]; and [redacted] response dated [redacted] to our customer questionnaire dated [redacted], question [redacted].

being a place for customers to learn about a product, evaluate its appropriateness for their needs, and access ongoing product support.²⁰¹ Hyperscalers' websites provide a clear demonstration of this in practice, with an extensive range of information, guidance, and support being offered by each.

- 4.61 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.
- 4.62 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 5.
- 4.63 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

- 4.64 Each of the hyperscalers operate co-sell schemes.²⁰² Co-selling schemes can include a variety of features such as creating visibility for the ISVs solutions with the hyperscalers' sales teams 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.²⁰³ Indeed, this can even include apps which already combine

²⁰¹ <https://aws.amazon.com/marketplace/campaigns/dcx/chapter3/>, <https://learn.microsoft.com/en-us/marketplace/azure-marketplace-overview>, <https://cloud.google.com/marketplace/docs>.

²⁰² 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 7 February 2023]. Google website. [Google Cloud doubles-down on ecosystem in 2022 to meet customer demand | Google Cloud Blog](#) [accessed 8 February 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 8 February 2023].

²⁰³ AWS website. [AWS ISV Accelerate Program \(amazon.com\)](#) [accessed 7 February 2023]. Microsoft website. [Marketplace rewards - your commercial marketplace benefits - Marketplace publisher | Microsoft Learn](#) [accessed 7 February 2023]. Google website. [Google Cloud doubles-down on ecosystem in 2022 to meet customer demand | Google Cloud Blog](#) [accessed 7 February 2023].

products from ISVs and the hyperscalers, as seen with AWS's partnership with Salesforce.²⁰⁴

- 4.65 The hyperscalers often tailor the level of support provided to ISVs to reflect the level of integration between the ISV and the cloud provider.²⁰⁵ Access to Microsoft tiers which offer higher levels of support requires tight technical integration with Azure and driving Azure-based sales revenue.²⁰⁶ We also understand that AWS has [redacted].²⁰⁷

Marketplaces

- 4.66 Each of the hyperscalers operates a cloud marketplace. A cloud marketplace is an online platform allowing cloud providers and ISVs to sell their services to business customers. Table 4.9 below sets out key data in relation to the UK marketplaces of each of AWS, Google and Microsoft.
- 4.67 [redacted] operates the largest marketplace [redacted]. ISVs transact significantly more [redacted] on [redacted] than on [redacted]. In addition, [redacted] and is [redacted].²⁰⁸
- 4.68 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 ([redacted]%) of gross sales revenue on its marketplace are third-party sales, with listing fee revenue from third-party sales accounting for [redacted]% of its marketplace revenues.²⁰⁹ Our market research indicated that AWS users are more likely to use marketplaces to buy third-party services than all marketplace users.²¹⁰ 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.²¹¹

²⁰⁴ AWS website. <https://aws.amazon.com/featured-partners/salesforce/> [accessed 7 February 2023].

²⁰⁵ [redacted] 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 [redacted] ISVs). See [redacted] response dated [redacted] to s.174 request of [redacted], [redacted]. [redacted] notes that “The intent [in 2022] is to have a consistent and fair set of partner benefits and compensation treatment for all ISVs [redacted] in the same tier to prioritize those partners that will have the highest business impact for customers.” See [redacted] response dated [redacted] to s.174 request of [redacted], [redacted].

²⁰⁶ 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 1 March 2023].

²⁰⁷ AWS response dated 28 November 2022 to s.174 request of 24 October 22, [redacted].

²⁰⁸ [redacted] response dated [redacted] to s.174 request of [redacted], [redacted]; and [redacted] response dated [redacted] to s.174 request of [redacted], [redacted].

²⁰⁹ Ofcom analysis of AWS response dated 9 December 2022 to s.174 request of 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.

²¹⁰ Context Consulting research data tables, Q46.

²¹¹ Microsoft response dated 9 December 2022 to s.174 request of 21 October 2022, question 19; and its response dated 16 January 2022 to follow-up email dated 20 December 2022. Context Consulting research data tables, Q46.

Table 4.9: Gross sales revenue and commission revenue received from third-party products through AWS, Azure and Google Marketplaces, in the UK (2021)

| | Gross sales revenue of first-party products | Gross sales revenue of third-party products | Commission revenue from sales of third-party products |
|---------------------------|---|---|---|
| AWS Marketplace | [redacted] | [redacted] | [redacted] |
| Azure Marketplace | [redacted] | [redacted] [£10-20m] | [redacted] [<£5m] |
| Google Marketplace | [redacted] | [redacted] | [redacted] |

Sources: Figures for AWS response dated 21 November 2022 to s.174 notice of 24 October 2022, question 21a-c; AWS response dated 17 February 2023 to follow-up email dated 28 January 2023 to s.174 notice of 24 October 2022, paragraph 14.1 and AWS response dated 9 December 2022 to s.174 notice of 24 October 2022, question 19, Annex 2. Figures for Microsoft: Microsoft response dated 16 January 2023 to follow-up email of 20 December 2022 to s.174 notice of 21 October 2022, and Microsoft response dated 9 December 2022 to s.174 notice dated 21 October 2022, question 19, Annex B5. For Google: Gross sales of third-party services data from Google response dated 22 December 2022 to s.174 request of 26 October 2022, question 19, page 39. Figures converted to GBP using average USD exchange rate for 2021. 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.²¹² Google and Microsoft were unable to provide the value of first-party sales via their marketplaces.²¹³ Gross sales refers to the total value of the sales transacted through the marketplace.

4.69 At present, marketplaces are not a major revenue source for the hyperscalers, accounting for a very small proportion of their total cloud revenues (c. [redacted]%).²¹⁴ However, there are some indicators that marketplaces may grow in importance over time. Some commentators expect that marketplaces will continue to grow.²¹⁵ Evidence suggests that the hyperscalers see marketplaces as an important distribution channel and are seeking to grow their use. [redacted] aims for its marketplace to be the primary online distribution channel

²¹² AWS response dated 21 November 2022 to s.174 notice of 24 October 2022, Q21a and AWS response dated 9 December 2022 to s.174 notice of 24 October 2022, question 19b.

²¹³ Microsoft response dated 9 December 2022 to s.174 request of 21 October 2022, question 19, and its response dated 16 January 2023 to follow-up email dated 20 December 2022; and Google response dated 22 December 2022 to s.174 request of 26 October 2022, question 19, page 39.

²¹⁴ Based on the revenue from first-party sales and commission fees from third-party sales. [redacted] response dated [redacted] to s.174 request of [redacted], question [redacted]; [redacted] response dated [redacted] to s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to s.174 request of [redacted], question [redacted]; [redacted] response dated [redacted] to s.174 notice of [redacted], question [redacted]; [redacted] response dated [redacted] to s.174 notice dated [redacted], question [redacted], and [redacted] response dated [redacted] to s.174 notice dated [redacted], question [redacted].

²¹⁵ For example, see Canalys Insights. [Canalys Insights - Are cloud marketplaces worth the hype?](#) [accessed 27 March 2023]. Also see [redacted] response dated [redacted] to s.174 request of [redacted], [redacted]; [redacted] response dated [redacted] to s.174 request of [redacted], [redacted].

for both first-party and third-party solutions.²¹⁶ [X] aims to become the most strategic channel that ISVs use to acquire new customers,²¹⁷ whereby [X].²¹⁸

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

4.70 Evidence gathered from the hyperscalers suggests that they, [X], 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. [X] explains that ISVs are "an important sell-through channel" which are "actively selling solutions that run on [X], helping us acquire new customers and driving more [X] consumption from existing customers".²¹⁹ [X] explained the critical importance of being a first-choice platform for ISVs as IT spend is shifting to software rather than infrastructure consumption.²²⁰ [X] explains that "marketplace is built primarily to drive underlying [X] consumption",²²¹ and "marketplace partners will be offered increasingly tiered incentives and programs [X] depending on how much they drive infrastructure use ([X]) and also [X] ([X])".²²²

Question 4.1 Do you agree with our assessment of how customers buy cloud infrastructure services and how cloud providers seek to acquire customers?

Market outcomes

4.71 In this subsection we report key market outcomes, including the positioning of the major cloud providers in the market. We look at the shares of supply for UK cloud services, analyse the profitability of the key providers, and consider the service offerings and key capabilities of each hyperscaler.

UK shares of supply

UK revenues for IaaS and PaaS

4.72 Table 4.10 summarises our estimates of UK revenues associated with IaaS and PaaS. We explain how we derived these estimates in Annex 5. We have focused on UK shares of supply rather than global shares, as global shares would likely understate the position of the hyperscalers in the UK, given the large Asian cloud providers, such as Alibaba, have a more limited presence in the UK.

²¹⁶ [X] response dated [X] to s.174 request of [X], [X].

²¹⁷ [X] response dated [X] to s.174 request of [X], [X].

²¹⁸ [X] response dated [X] to s.174 request of [X], [X].

²¹⁹ [X] response dated [X] to s.174 request of [X], [X].

²²⁰ [X] response dated [X] to s.174 request of [X], [X].

²²¹ [X] response dated [X] to s.174 request of [X], [X].

²²² [X] response dated [X] to s.174 request of [X], [X].

4.73 We estimate that in 2021, cloud infrastructure services generated revenues of £4.5bn to £5.0bn. Between 2019 and 2021, UK revenues for IaaS and PaaS combined grew by 25% - 30% per year.

Table 4.10: UK IaaS and PaaS revenues, £bn

| | 2019 | 2020 | 2021 | Annual growth |
|---------------|---------------|---------------|---------------|---------------|
| IaaS | [£] [1.5-2.0] | [£] [2.0-2.5] | [£] [2.5-3.0] | 25% - 30% |
| PaaS | [£] [0.5-1.0] | [£] [1.0-1.5] | [£] [1.5-2.0] | 35% - 40% |
| IaaS and PaaS | [£] [2.5-3.0] | [£] [3.5-4.0] | [£] [4.5-5.0] | 25% - 30% |

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 2021.

4.74 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 5.

4.75 Based on data obtained from International Data Corporation (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.²²³

4.76 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 2021.²²⁴

4.77 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.

4.78 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 2024/2025,²²⁵ though recent earnings releases from the largest providers indicate global revenue growth could be slower in 2023.²²⁶

²²³ IDC, 2022. Public Cloud Services Tracker, 2021 H2 (published April 2022).

²²⁴ Annex 5, Shares of supply for cloud infrastructure services in the UK.

²²⁵ [£].

²²⁶ For example, in its Q4 2022 earnings call, Amazon said AWS's year on year revenue growth had slowed as enterprises optimised cloud spending in response to macroeconomic conditions, and revenue growth in the first month of 2023 was in the mid-teens – down from 20% growth in Q4 2022. [Amazon Q4 2022 earnings call](#) (6 minutes 50 seconds to 7 minutes 50 seconds).

UK shares of supply for IaaS

- 4.79 The table below shows our estimated shares of supply for UK IaaS. We explain how we derived these estimates in Annex 5.
- 4.80 Based on responses to our information requests, and the market size estimates set out above, we estimate AWS and Microsoft accounted for approximately [X] [70% to 80%] of UK IaaS revenues in 2021; a proportion that has remained broadly stable since 2019. Within this, Microsoft’s share of UK IaaS revenues has grown slightly while AWS’s share has reduced slightly.

Table 4.11: UK IaaS shares of supply, 2019 – 2021

| | 2019 | 2020 | 2021 |
|------------------------|--------------|--------------|--------------|
| AWS | [X] [40-50%] | [X] [40-50%] | [X] [40-50%] |
| Microsoft | [X] [30-40%] | [X] [30-40%] | [X] [30-40%] |
| AWS + Microsoft | [X] [70-80%] | [X] [70-80%] | [X] [70-80%] |
| Google | [X] [0-5%] | [X] [0-5%] | [X] [0-5%] |
| Other | [X] [10-20%] | [X] [10-20%] | [X] [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.

- 4.81 While Google’s UK IaaS revenues have grown since 2019, it represented [X] [0% to 5%] of UK IaaS revenues in 2021, significantly behind AWS and Microsoft.
- 4.82 The ‘other’ category includes a number of other providers with low IaaS shares of supply in the UK, which have steadily declined as Microsoft and Google have gained share. This category includes IBM and Oracle, though their shares were both [X] [0% to 5%] in 2021, around half that of Google. While UK IaaS revenues for IBM and Oracle grew between 2019 and 2021, [X].

UK shares of supply for PaaS

- 4.83 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.²²⁷ The table below shows our estimated shares of supply for UK PaaS. We explain how we derived these estimates in Annex 5.

²²⁷ Based on IDC’s categorisations – see Annex 5.

Table 4.12: UK PaaS shares of supply, 2019 – 2021

| | 2019 | 2020 | 2021 |
|------------------------|--------------|--------------|--------------|
| AWS | [X] [20-30%] | [X] [20-30%] | [X] [20-30%] |
| Microsoft | [X] [10-20%] | [X] [20-30%] | [X] [20-30%] |
| AWS + Microsoft | [X] [40-50%] | [X] [40-50%] | [X] [40-50%] |
| Google | [X] [5-10%] | [X] [10-20%] | [X] [10-20%] |
| Other | [X] [40-50%] | [X] [40-50%] | [X] [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.

- 4.84 Table 4.12 indicates that AWS and Microsoft represented [X] [40% to 50%] of UK PaaS revenues in 2021 – a lower share than for IaaS. Within this, we estimate that Microsoft’s share of UK PaaS revenues has grown slightly while AWS’s share has remained stable.
- 4.85 Google’s share is closer to that of AWS and Microsoft in PaaS than in IaaS, with a [X] [10% to 20%] share of UK PaaS revenues in 2021. [X].
- 4.86 Oracle and MongoDB have some of the larger shares of supply of companies in the ‘other’ category, followed by IBM; we estimate they represented around [X] [0% to 5%], [X] [0% to 5%] and [X] [0% to 5%] respectively of UK PaaS revenues in 2021.

UK shares of supply for IaaS and PaaS combined

- 4.87 Table 4.13 shows our estimated shares of supply for UK IaaS and PaaS combined, drawing on the information presented above.

Table 4.13: UK shares of supply for IaaS and PaaS combined, 2019 – 2021

| | 2019 | 2020 | 2021 |
|------------------------|--------------|--------------|--------------|
| AWS | [X] [30-40%] | [X] [30-40%] | [X] [30-40%] |
| Microsoft | [X] [20-30%] | [X] [30-40%] | [X] [30-40%] |
| AWS + Microsoft | [X] [60-70%] | [X] [60-70%] | [X] [60-70%] |
| Google | [X] [0-5%] | [X] [5-10%] | [X] [5-10%] |
| Other | [X] [20-30%] | [X] [20-30%] | [X] [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.

- 4.88 We estimate that in 2021 AWS and Microsoft had around [X] [60% to 70%] share of UK combined IaaS and PaaS revenues, with Google significantly lower on [X] [5% to 10%].

Overall, we estimate that AWS, Microsoft and Google accounted for [§<] [70% to 80] of UK IaaS and PaaS revenues in 2021.

- 4.89 Between 2019 and 2021, the hyperscalers' overall UK share of supply increased. Over this period, we estimate that AWS's UK share of supply fell marginally, Microsoft's grew slightly and Google experienced stronger growth, although from a lower revenue base.
- 4.90 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 2021.

UK shares of supply for SaaS

- 4.91 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.
- 4.92 In 2021, the hyperscalers' share of UK 'SaaS – Applications' revenue was around 13% according to IDC, most of which related to Microsoft services.²²⁸

Positioning of the hyperscalers

- 4.93 In this subsection, we describe in more detail the positioning of AWS, Microsoft and Google, as the leading players in the cloud infrastructure services market 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

- 4.94 AWS is generally recognised as offering a strong portfolio of services across both IaaS and PaaS. In 2021, AWS's public cloud infrastructure revenues were primarily driven by its [§<] and [§<] product categories.²²⁹ 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.²³⁰
- 4.95 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.²³¹

²²⁸ IDC, 2022. Public Cloud Services Tracker, 2021 H2 (published April 2022).

²²⁹ AWS response dated 9 December 2022 to s.174 request of 24 October 2022, Part B, question 4.

²³⁰ Computer Weekly, 2019. [Should AWS consider building out its SaaS play?](#) [accessed 9 February 2023].

²³¹ [§<] response dated [§<] to the s.174 notice dated [§<], question [§<]; and [§<] response dated [§<] to the s.174 notice dated [§<], question [§<].

4.96 AWS today is recognised to have the greatest breadth and depth of capabilities,²³² due to the range of services it offers, and the functionality offered within those services.²³³ As noted above, AWS also [redacted]. One source [redacted] suggests that AWS positions itself as the best provider of cost effective, scalable infrastructure and platform as a service,²³⁴ with another source suggesting that AWS may be particularly cheap for compute infrastructure due to its Graviton processor.²³⁵ We also observe that BT, [redacted] and [redacted] 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).²³⁶

Microsoft

- 4.97 Microsoft is generally recognised to have strong capabilities across both IaaS and PaaS. In 2021, Azure's revenue was mainly driven by its compute and storage product categories.²³⁷ However, customers also recognise Microsoft and/or may choose Microsoft for its enterprise SaaS offerings.²³⁸
- 4.98 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.²³⁹ 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.²⁴⁰
- 4.99 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.²⁴¹ 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.²⁴² This may be particularly important for customers that require highly integrated features across their on-premises and cloud infrastructures (e.g. banks).

²³² [redacted].

²³³ This includes its own cloud services, as well as third-party services (e.g. from ISVs).

²³⁴ [redacted] response dated [redacted] to our proposed use of information dated [redacted].

²³⁵ Redburn, 2022. Cloud Computing report, page 12. The Graviton processor is used to power Amazon EC2 instance types. [What is AWS Graviton?](#) [accessed 22 March 2023].

²³⁶ [BT response](#) to the CFI, page 19; [redacted] response to the CFI, page [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

²³⁷ Microsoft response dated 12 December 2022 to s.174 request of 21 October 2022, Part B, questions 1 & 4.

²³⁸ [BT response](#) to the CFI, page 19; and [redacted] response dated [redacted] to our customer questionnaire, question [redacted].

²³⁹ Datamation, 2023. [AWS vs. Azure vs. Google Cloud](#) [accessed 7 March 2023].

²⁴⁰ Context Consulting research data tables, Q25.

²⁴¹ Context Consulting research report, Slide 57.

²⁴² Microsoft website. [Azure Active Directory](#) [accessed 17 February 2023].

Google

- 4.100 In 2021, Google Cloud's revenues were primarily driven by its [redacted] and [redacted] product categories.²⁴³ 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.²⁴⁴
- 4.101 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.²⁴⁵ This specialism is recognised by a range of cloud providers and customers alike. We also note that analytics is [redacted].²⁴⁶
- 4.102 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,²⁴⁷ having been responsible for first introducing Kubernetes, which enables the management of containerised applications. We observe that [redacted].²⁴⁸
- 4.103 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).²⁴⁹ 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.²⁵⁰

Profitability

- 4.104 In this subsection, we:
- compare hyperscaler operating profits and margins to those of other cloud providers; and
 - compare our estimates of return on capital employed (ROCE) for AWS and Microsoft's cloud businesses to the weighted average cost of capital (WACC).²⁵¹

²⁴³ Google response dated 16 December 2022 to s.174 request of 26 October 2022, Part B, questions 1 & 4, Annex 2.

²⁴⁴ Google website. [Google Workspace](#) [accessed 17 February 2023].

²⁴⁵ BigQuery is a fully managed enterprise data warehouse with built-in features like machine learning, etc.

²⁴⁶ Ofcom analysis of [redacted] response dated [redacted] to s.174 request of [redacted], [redacted]; [redacted] response dated [redacted] to s.174 request of [redacted], [redacted]; and [redacted] response dated [redacted] to s.174 request of [redacted], [redacted].

²⁴⁷ TechCrunch, 2019. [Google remains the top open-source contributor to CNCF projects](#) [accessed 7 March 2023].

²⁴⁸ [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted]; and [redacted] response to [redacted] of our follow-up email dated [redacted] concerning the s.174 notice dated [redacted].

²⁴⁹ Google's response dated 23 November 2022 to the s.174 notice dated 26 October 2022, question 8.

²⁵⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

²⁵¹ The definition of ROCE and WACC can be found in Annex 8 (Glossary).

- 4.105 Our profitability analysis focuses on AWS, Azure and Google Cloud as the hyperscaler businesses providing cloud infrastructure services.²⁵² 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 365.²⁵³ 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.
- 4.106 We have run our 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.
- 4.107 We explain our approach to assessing profitability in more detail in Annex 6.

Hyperscaler operating profits compared to other global cloud providers

- 4.108 Amazon and Google publicly report earnings before interest and tax (EBIT) for AWS and Google Cloud respectively.
- 4.109 Microsoft does not publicly report EBIT for Azure or Microsoft Cloud.²⁵⁴ 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 6.
- 4.110 Figure 4.14 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.²⁵⁵ IBM and Oracle are not included in the chart as IBM provided only high-level information relating to cloud EBIT, and Oracle said it could not provide detailed or accurate estimates on cloud profit beyond its public reporting for the Cloud and License segment.²⁵⁶
- 4.111 Figure 4.14 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

²⁵² 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 6.

²⁵³ Microsoft, 2022. [2022 10-K report](#), page 42 [accessed 23 March 2023].

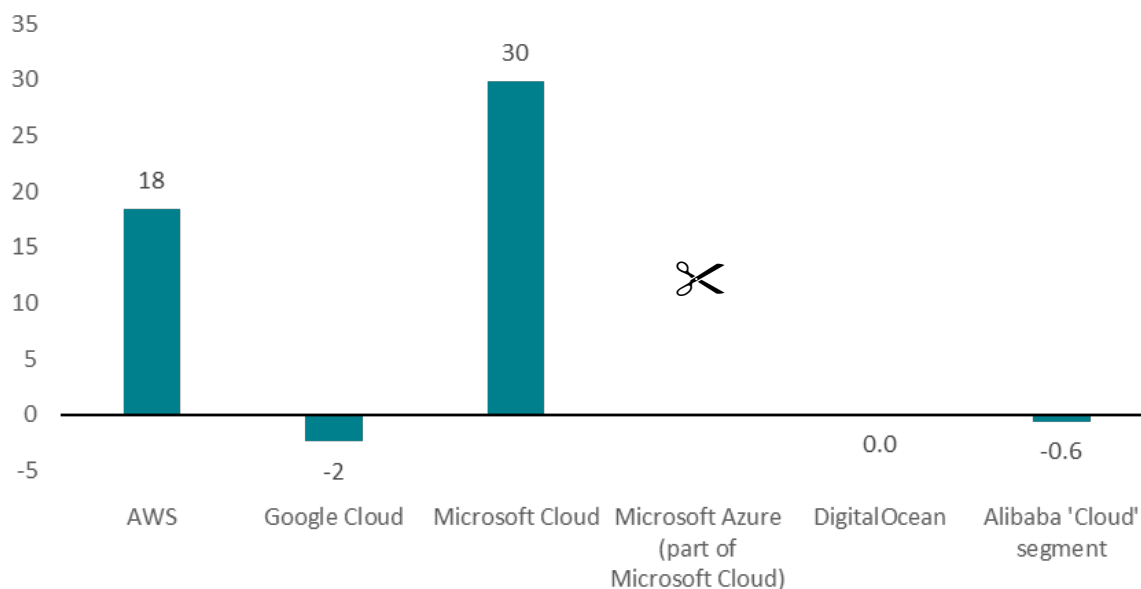
²⁵⁴ While Microsoft does publish EBIT for its Intelligent Cloud operating segment, this includes Azure alongside non-cloud server products, so it may not be representative of the profits associated with cloud infrastructure services.

²⁵⁵ DigitalOcean provides cloud infrastructure services for startups and small and medium-sized businesses.

²⁵⁶ 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.

providers, including Google Cloud, which is loss making.²⁵⁷ We estimate that Azure EBIT in Microsoft’s 2022 financial year was [redacted]. This is [redacted].

Figure 4.14: Annual global EBIT for the latest financial year (£bn)



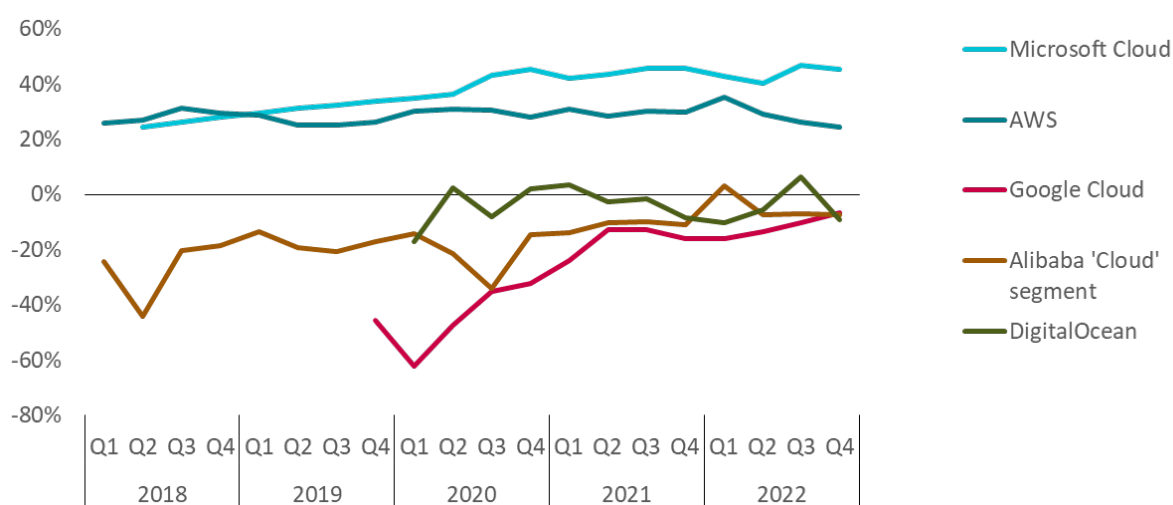
Source: Ofcom analysis of cloud providers’ published financial statements, information provided by Microsoft in response to our information requests²⁵⁸ and Ofcom assumptions.

4.112 Figure 4.15 shows quarterly EBIT margins for the same businesses between 2018 and 2022. Our estimates for Azure have been redacted in the published version of this report.

²⁵⁷ 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.

²⁵⁸ 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.

Figure 4.15: Quarterly EBIT margins for major cloud providers, 2018-22



Source: Ofcom analysis of cloud providers’ financial data reported by S&P Capital IQ, information provided by Microsoft in response to our information requests²⁵⁹ and Ofcom assumptions. We have used the latest filings (incorporating restatements).

4.113 Figure 4.15 indicates that, over this period:

- AWS had stable EBIT margins of around 20% to 30%
- 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 [30].
- Google Cloud was loss making despite strong revenue growth, though quarterly losses have steadily reduced.
- Among other cloud providers, quarterly EBIT margins for Alibaba’s ‘Cloud’ segment and DigitalOcean were occasionally positive but mostly negative in recent years.

4.114 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 [30].

Comparison of hyperscaler returns against WACC

4.115 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.²⁶⁰ When combined with other indicators, a finding that ROCE

²⁵⁹ 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.

²⁶⁰ ROCE is calculated by dividing EBIT by the value of capital employed in the relevant business.

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 6.

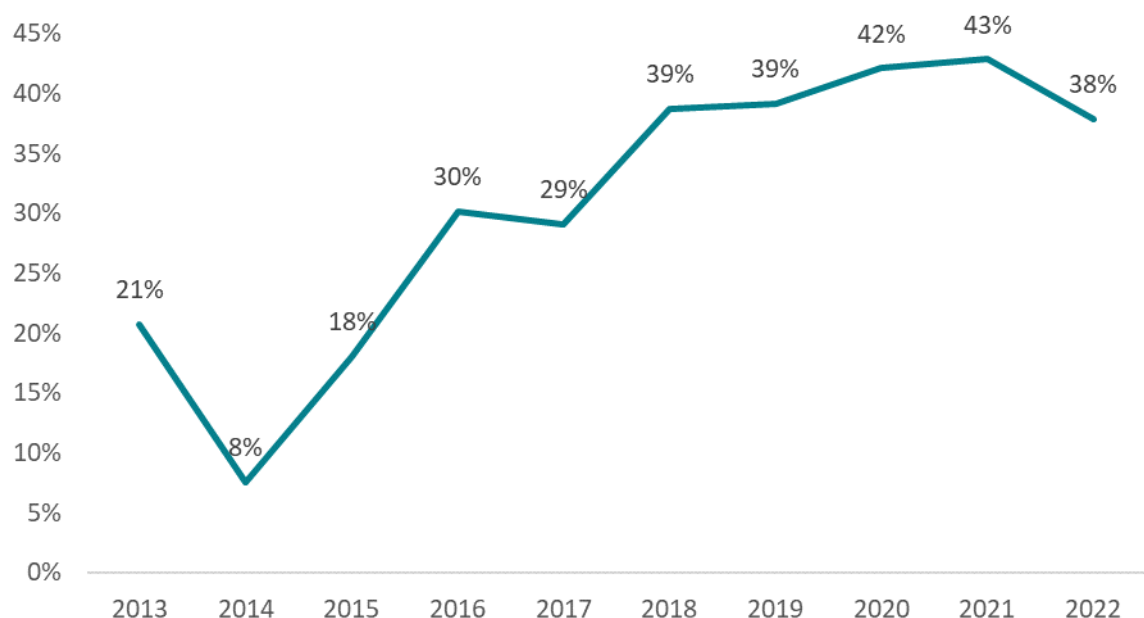
4.116 Our ROCE analysis focuses on AWS and Microsoft Azure, as our shares of supply analysis indicates they represent a substantial share of cloud infrastructure revenues, and we know that Google Cloud is currently loss-making and will have a negative ROCE. For context and comparison with Azure, we also present an estimate of ROCE for Microsoft Cloud.

4.117 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 6.

AWS

4.118 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 2013.²⁶¹ Our estimates are shown in the figure below.

Figure 4.16: AWS ROCE estimates, 2013-22



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

4.119 Our estimate indicates that AWS ROCE has increased since 2013 and has been around 40% since 2018. It is higher than our estimate of WACC in all years except 2014. In Annex 6 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.

²⁶¹ A 10-K form is an annual report required by the Securities and Exchange Commission in the US. It includes annual financial statements.

Microsoft

- 4.120 Annex 6 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.
- 4.121 Our ROCE estimates for Azure for Microsoft’s financial years ending 2018 to 2022 are shown in the figure below.
- 4.122 We estimate that Azure’s ROCE increased over this period and is now above our estimate of WACC. [X].

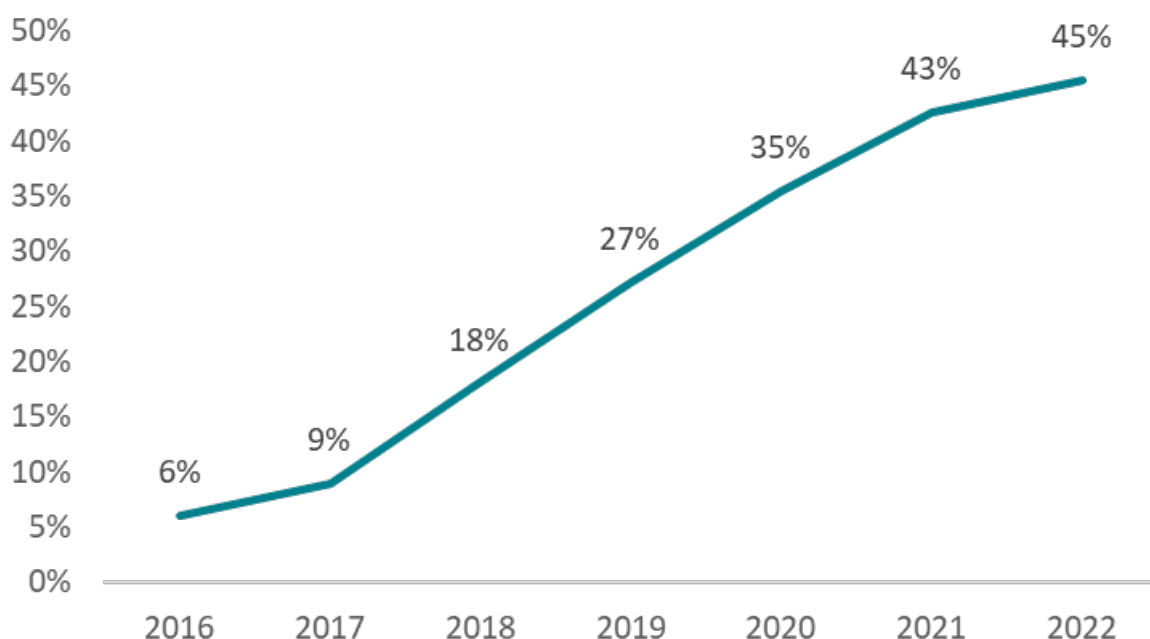
Figure 4.17: Azure ROCE estimate, Microsoft financial years ending 2018 - 2022

[X]

Source: Ofcom analysis based on Microsoft 10-K reports and information provided by Microsoft in response to our information requests.²⁶²

- 4.123 Our baseline ROCE estimates for Microsoft Cloud for Microsoft’s financial years ending 2016 to 2022 are shown in the figure below.²⁶³

Figure 4.18: Baseline ROCE estimate for Microsoft Cloud, Microsoft financial years ending 2016 to 2022



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

²⁶² 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.

²⁶³ As noted in Annex 6, 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.

- 4.124 We estimate that Microsoft Cloud ROCE steadily increased over this period. Our baseline estimate of Microsoft Cloud ROCE increased from 18% to 45% in Microsoft's financial years ending 2018 to 2022, above our estimate of WACC.
- 4.125 In Annex 6 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.

Question 4.2: Do you agree with our characterisation of the market outcomes in supply of cloud infrastructure services?

5. Potential barriers to effective competition

- 5.1 In Section 4 we explained, on a preliminary basis, that cloud infrastructure is relatively concentrated in the UK with AWS and Microsoft accounting for the majority of revenues across IaaS and PaaS. In this section we consider the extent to which the market outcomes we describe in Section 4 are determined by significant barriers to effective competition in the supply of public cloud infrastructure services.
- 5.2 In doing so, we have not carried out a formal market definition assessment, but instead considered the extent to which different market features and business practices may inhibit effective competition.
- 5.3 Unless otherwise stated, our assessment applies to all layers of cloud infrastructure services (i.e. across IaaS and PaaS). This is because customers often buy groups of cloud infrastructure services across different layers of the cloud stack.
- 5.4 This section is split into three parts:
- a) First, we consider the extent to which cloud providers are constrained by customers switching or multi-clouding. In doing so, we assess the nature and scope of different barriers to switching and multi-cloud.
 - b) Second, we consider the extent to which AWS and Microsoft are constrained by the threat of entry or expansion by competing cloud providers, including Google. In doing so, we assess the nature and scope of different barriers to entry and expansion.
 - c) Third, we consider the extent to which the hyperscalers may be using their roles as suppliers to ISVs and distributors of ISVs cloud services to engage in practices that may increase barriers to entry and expansion for ISVs.

Barriers to switching and multi-cloud

- 5.5 In this subsection we consider the extent to which barriers to switching and multi-cloud may weaken effective competition and adversely affect market outcomes.
- 5.6 The evidence received indicates that some customers may face three categories of barriers to switching and multi-cloud:
- a) **Technical barriers.** These include: (i) the technical efforts required to set-up a preferred cloud architecture and tailor applications so they can run on it; and (ii) the 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 (in particular the hyperscalers) 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.

5.7 In the following paragraphs we consider the extent to which each of these barriers may limit customers' ability to switch or multi-cloud and assess whether these obstacles may be particularly strong for specific use-cases or customer segments.

Lack of technical interoperability and portability

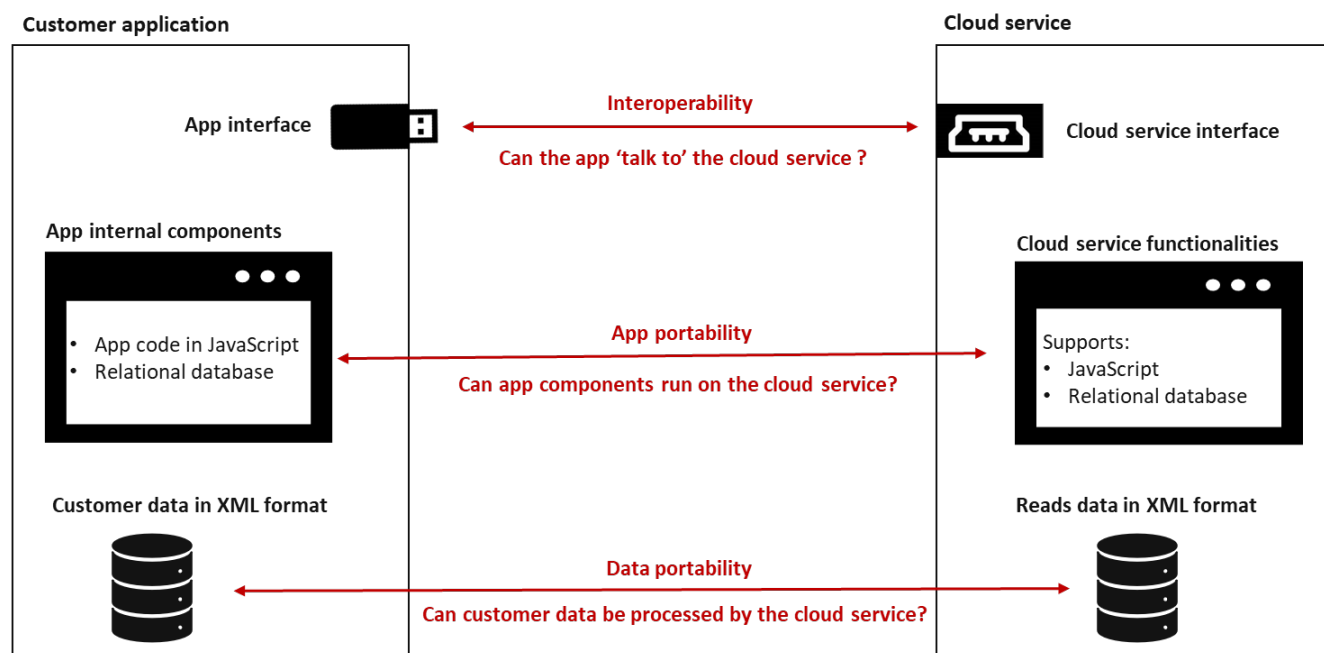
5.8 In this subsection we examine the extent to which a lack of interoperability and portability may hinder customers' ability to switch and multi-cloud. The subsection is structured as follows:

- a) First, we introduce the concepts of interoperability and portability in cloud and outline how they may affect customers' ability to switch and multi-cloud.
- b) Second, we discuss the evidence received from the hyperscalers and customers with a view to setting out our preliminary assessment of the extent to which a lack of interoperability and portability may hinder the ability of cloud customers to implement different multi-cloud architectures and switch.
- c) Finally, we present evidence suggesting that technical barriers in cloud may be exacerbated by specific business practices.

A lack of interoperability and portability can hinder customers' ability to switch and multi-cloud

5.9 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 applications and data work and perform the required tasks in the cloud, customers typically face three categories of technical challenges: interoperability, application portability and data portability challenges. These challenges are discussed below in relation to a hypothetical example of a customer running a data analytics application on a compute resource offered by a cloud provider.

Figure 5.1: Stylised example of customer application relying on a cloud service



Source: Ofcom.

5.10 As depicted in the above figure, customers wanting to run their analytics application on a cloud face three categories of technical challenges:

- a) **Interoperability:** each cloud provider typically defines and publishes a set of rules (e.g. APIs, protocols and workflows) that customers must follow to enable their applications to exchange data and information (i.e. interoperate) with a storage service on their cloud. Interoperability is a key consideration if a customer wants to run their application on the cloud of a different or additional cloud provider. To the extent the interoperability rules are different across cloud providers, the customer has to change some parts of their application so it can ‘talk to’ the storage service offered by the different/additional cloud provider.²⁶⁴ Similarly, interoperability is a key consideration if a customer wants to integrate different cloud services hosted on several clouds.²⁶⁵ To the extent the two cloud providers use incompatible interoperability rules, the customer may need to develop or procure an adaptor²⁶⁶ to connect their cloud services.
- b) **Application portability:** customers who architect and engineer their software for the cloud typically own their applications and data. However, they outsource operations to

²⁶⁴ For example, a customer may need specific code to read data from storage on cloud A and different code for performing the same action on cloud B. To use the analogy depicted in the figure, if cloud B uses a HDMI port (as opposed to a USB port) the customer needs to change the port of their service or use an adaptor when switching.

²⁶⁵ For example, suppose that in addition to the storage service a customer wants to integrate with an analytics service hosted on a different cloud.

²⁶⁶ An adaptor is a piece of software that intermediates the communication between two or more components that cannot directly interoperate with each other. Adaptors can act as abstraction layers that translate and bridge communication between otherwise incompatible APIs, formats and protocols.

cloud infrastructure services. To port an application to a different or additional cloud, the target cloud needs to support the application, the data and provide equivalent service capabilities to satisfy all dependencies. If not, some re-engineering work may be needed. For example, suppose a customer who is running an application as a cloud function on cloud A wants to run their application as a cloud function on cloud B. If cloud B only supports cloud functions written in the programming language JavaScript (and not Python), the application code needs to be translated from Python to JavaScript before the application could run on it.²⁶⁷

- c) **Data portability:** when using cloud infrastructure services, customers are typically in control of the data and how it is organised, with the cloud service offering basic storage capabilities such as a file system or object store. This is a relevant consideration when a customer wants to store their data in a different or additional cloud, because the target cloud needs to support the same storage capabilities. However, as most clouds offer basic storage capabilities (e.g. they can all store and process data in XML format), data portability is typically less of an issue compared to interoperability and application portability.²⁶⁸

- 5.11 Overall, a lack of interoperability can be regarded as a high degree of differentiation between interfaces of different cloud services ('can they all be accessed in the same manner?'), whereas a lack of portability can be regarded as a high degree of differentiation between functionalities of different cloud services ('do they all support a particular technology?'). The greater the degree of technical differentiation across any of these dimensions, the more effort is required of customers to run their applications on different or additional clouds.
- 5.12 On the contrary, the use of open-source or commercial cloud services that are available across clouds (i.e. cloud-agnostic services) can standardise the set of functionalities offered by different clouds and therefore mitigate portability challenges. Similarly, the use of standardised rules (e.g. open-APIs and open protocols) for interoperating with comparable cloud services can standardise interfaces and thus mitigate interoperability concerns.

Hyperscalers said they provide services to facilitate switching and multi-cloud but our analysis suggests take up of these services may be limited

- 5.13 Each of the hyperscalers submitted that they offer a wide range of support for switching and multi-cloud, including open-source software and standards, cloud services and tools designed to facilitate switching and multi-cloud and advice to customers. In particular:
 - a) AWS said its customers have complete freedom to run their applications anywhere and to reach any end user. AWS said its customers are also free to adopt multi-cloud

²⁶⁷ Note that application portability is a separate consideration from interoperability because, even if the customer could launch comparable cloud functions on cloud B, their application would still be unable to run there, if the necessary technologies are not supported by the target cloud.

²⁶⁸ In some cases, customer applications may rely on specific data storage solutions that may not or only partially exist in the target cloud. In those cases, data portability may become more complex, but it remains generally less problematic than interoperability and application portability.

strategies to meet their varying IT needs, without being bound to a specific cloud provider. To support this, AWS said it often builds its services on, or using, open-source technologies and standards.²⁶⁹ AWS said it makes a wide range of services and tools available to customers who wish to migrate,²⁷⁰ or multi-cloud.²⁷¹ AWS also said it educates its customers on building for “reversibility” in their IT solutions – i.e. the ability to retrieve and move data from one IT environment to another – including by advising on application design strategies to maximise portability, and encourages its customers to use open standard software and components to allow for multi-cloud solutions.²⁷² Finally, AWS also said it allows third parties to use AWS APIs and software development kits (SDKs) outside AWS.^{273, 274}

- b) Microsoft said that switching costs between cloud providers depends on the nature of the solution and are largely in the customer’s control. Microsoft said customers of cloud services are typically sophisticated organisations which enables them to make well-informed choices with a clear understanding of the trade-offs involved in optimising workloads to reduce the costs of migration. It said this may involve customers using technical solutions to mix and match cloud services from different cloud providers.²⁷⁵ Microsoft said that open-source technologies (e.g. Linux, Kubernetes) facilitate switching because they are prevalent across public clouds. It also said that 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. Microsoft said it offers customers help with

²⁶⁹ AWS said it supports: open standards (e.g. Linux operating systems and SQL database engines); REST and HTTP APIs that support OIDC and OAuth 2.0 authorisation; internet security protocols (e.g. SSL/TLS) and ciphers (e.g. RSA ciphers), identity protocols (e.g. SAML 2.0 and SCIM); devices and clients that use the MQTT and the MQTT over WebSocket Secure communications protocols; devices and clients that use the HTTPS protocol; and service configuration using JSON and YAML files, X.509 SSL/TLS certificates (AWS response dated 31 October 2022 to the s.174 notice dated 24 October 2022, pages 2-3; and AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, part A question 31).

²⁷⁰ AWS said its migration services and tools include S3/EC2 Data Transfer, Amazon S3 Transfer Acceleration, AWS Transfer for SFTP, AWS Direct Connect, AWS DataSync, AWS Outpost, AWS Application Migration Service, VMware Cloud on AWS, AWS Snow Family (Snowball/Snowmobile/Snowcone), AWS SCT, AWS Database Migration Service, Babelfish for Aurora PostgreSQL, and AWS MAP (AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, part A questions 11 and 17).

²⁷¹ AWS said its integration services and tools include Amazon EKS and ECS Anywhere, AWS Single Sign On, AWS CloudFormation and AWS Config (AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, part A question 31).

²⁷² AWS response dated 31 October 2022 to the s.174 notice dated 24 October 2022, pages 2-3; and AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, part A questions 11 and 17.

²⁷³ AWS said its approach has enabled many cloud providers to copy its Amazon S3 APIs and customers to use its AWS Encryption SDK at no cost (AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, part A question 31).

²⁷⁴ AWS response dated 31 October 2022 to the s.174 notice dated 24 October 2022, pages 2-3; and 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.

²⁷⁵ By way of example, Microsoft said that Walmart recently built a cloud native platform for its ecommerce retail website by leveraging Kubernetes (container technology) to enable its platform to run across two public clouds – Microsoft Azure and Google – and to take advantage of its own OpenStack-based cloud infrastructure (Microsoft response to the CFI, page 9). However, Walmart’s choice of cloud architecture may be driven by the fact that it competes directly with AWS in the retail e-commerce space. For example, see <https://www.bloomberg.com/news/newsletters/2022-07-12/walmart-cloud-weans-itself-off-of-microsoft-azure-google-cloud>.

switching.²⁷⁶ Finally, Microsoft said it makes extensive information available to developers about the services available in Azure and how customers, partners, or competitors can access that functionality.^{277, 278}

- c) Google said that traditional IT is characterised by restrictive licensing, closed ecosystems, and tying, whereas the cloud's original promise and potential is to be open, elastic, and free from artificial lock-ins. Google said it has an open source and open access approach to cloud. Specifically, Google said it is committed to using open APIs across its technology stack, building many of its services on open-source solutions, and giving customers options and tools to build, migrate and deploy their applications across multiple cloud environments to avoid vendor lock-in.^{279, 280, 281}

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

- a) The number of active customers using [X] increased from [X] in 2019 to [X] in 2021. These represent only around [X]% of the number of active customers using [X] in the same years.²⁸²
- b) The number of active customers using [X] increased from [X] in 2020 to [X] in 2022. This represents only around [X]% of the number of active customers using [X] in the same years.²⁸³

5.15 We recognise that the figures on number of active customers may not precisely reflect the individual number of users.²⁸⁴ However, we consider this analysis provides a good

²⁷⁶ Microsoft provided the examples of its Azure Migration and Modernization programme (Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, part A question 11).

²⁷⁷ Microsoft response to the CFI, pages 7-10; and Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, part A question 31.

²⁷⁸ Microsoft response to the CFI, pages 7-10 and footnote 16; and Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, part A questions 11 and 31.

²⁷⁹ For example, Google's Cloud Migration Center provides an end-to-end service for facilitating migration (see <https://cloud.google.com/migration-center/docs>). Google also offers more specialised migration support for services such as BigQuery and databases (see <https://cloud.google.com/bigquery/docs/dts-introduction> and <https://cloud.google.com/solutions/database-migration>).

²⁸⁰ Google said it was the first to launch a multi-cloud infrastructure service and the first to launch a multi-cloud data warehouse (Google response to the CFI, paragraphs 6, 15). Google also said it supports: open-source software including Linux and Kubernetes; cloud services focussed on interoperability including Google Anthos and Google BigQuery Omni; public APIs including Google Cloud APIs; and standard protocols including OAuth (Google response dated 23 November 2022 to the s.174 notice dated 26 October 2022, part A questions 31, 32, pages 59 and 60).

²⁸¹ Google response to the CFI, paragraphs 13, 14, 23 and 28; and Google response dated 23 November 2022 to the s.174 notice dated 26 October 2022, part A questions 2, 8, 17 and 31.

²⁸² To carry out this analysis we have compared the total number of active customers using [X] in 2019 and 2020 to the total number of active customers using [X] in the same years.

²⁸³ To carry out this analysis we have compared the total number of active customers using [X] in 2019 and 2020 to the total number of active customers using [X] in the same years.

²⁸⁴ [X] 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.

indication that take-up of some of the hyperscalers' services that they say are designed to facilitate switching and multi-cloud is likely to be limited.

5.16 In addition, [redacted].²⁸⁵

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

5.17 The market research found that 52% of customers cited a lack of interoperability between different IaaS/PaaS providers' services²⁸⁶ as a key concern with the cloud infrastructure services market.²⁸⁷ 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 switched (60%) and those who have added an IaaS/PaaS provider (56%). 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.

5.18 With regards to multi-cloud, at least half of IaaS/PaaS users in the market research only use a single cloud provider.²⁸⁸ This is likely due to technical barriers, at least in part. Our 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%).²⁸⁹ 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.

5.19 Similar to multi-cloud, 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 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

²⁸⁵ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

²⁸⁶ 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 7.

²⁸⁷ Context Consulting research data tables, Q63. The figure is 51% after excluding respondents using only private cloud.

²⁸⁸ 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 7.

²⁸⁹ 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%).

IaaS/PaaS provider was time and cost of making the change (43%).²⁹⁰ 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 than technically difficult (i.e. they know how to reconfigure their applications, but it would take time and money).

Multi-cloud: the evidence suggests that a lack of interoperability and portability can materially increase barriers to multi-cloud

5.20 This subsection considers the extent to which a lack of interoperability and portability may hinder customers’ ability to implement the following multi-cloud architectures:²⁹¹

- a) **Integrated multi-cloud:** the customer integrates different customer applications, customer data, and/or cloud services hosted on two or more public clouds into a consolidated architecture.
- b) **Cloud duplication:** the customer aims to mirror their cloud architecture on two or more public clouds so that all or some of their applications and data can work equivalently on all of them.
- c) **Siloed multi-cloud:** 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.

5.21 Overall, the evidence we have received indicates that using multiple clouds is always likely to require some technical effort, but its scale varies depending on the multi-cloud implementation. This is summarised in the table below and discussed in detail in the following paragraphs.

Table 5.2: Overview of preliminary findings on the technical effort to implement multi-cloud

| Multi-cloud scenario | Technical barriers | Customers most affected by barriers | Level of effort |
|----------------------------------|--|--|-----------------|
| a) Integrated multi-cloud | Interoperability effort to integrate separate public clouds. | Customers wishing to integrate a large number of different applications, data sets and/or cloud services hosted on separate public clouds. | High |

²⁹⁰ Context Consulting research data tables, Q52 and Q53. This figure is 43% after excluding respondents using only private cloud. This was also the first of the main barriers to completely switching the IaaS/PaaS provider (20%). This figure is 19% and maintains its relative ranking after excluding respondents using only private cloud.

²⁹¹ Our analysis of the different types of multi-cloud is based on stylised illustrations designed to capture the fundamentals for the purpose of our competition assessment and not technical descriptions of multi cloud architectures which would be much more complex.

| Multi-cloud scenario | Technical barriers | Customers most affected by barriers | Level of effort |
|------------------------------|--|--|-----------------|
| b) Cloud duplication | Interoperability and portability effort to reconfigure applications and data so they work equivalently on separate public clouds. Costs of maintaining parallel clouds. | Customers wishing to duplicate many workloads. Customers with a cloud architecture tightly integrated with proprietary services of cloud providers. | Medium to High |
| c) Siloed multi-cloud | Costs of maintaining parallel clouds. | Customers with a cloud architecture tightly integrated with proprietary services of cloud providers. | Medium |

Source: Ofcom.

(a) Interoperability and portability considerations when using multiple integrated clouds

- 5.22 With integrated multi-cloud different customer applications, customer data and/or cloud services hosted on two or more public clouds are highly integrated.
- 5.23 Integrated multi-cloud does not involve any porting of applications between clouds. However, for an integrated multi-cloud deployment to work effectively, different clouds need to exchange data in order to provide the functionalities required by the customer’s applications (e.g. Google BigQuery should be able to load data from Azure Storage). Hence, the key technical considerations revolve around the degree to which different customer applications, customer data and cloud services hosted on different clouds can interoperate with each other.
- 5.24 To the extent that different cloud providers use different interoperability rules (e.g. proprietary APIs, protocols and workflows), substantial technical effort is likely to be required to ensure that different clouds are able to ‘talk to each other’ (i.e. interoperate).²⁹² With respect to cloud services, even where they can ‘talk to each other’, services from different cloud providers may not be as tightly integrated as two comparable cloud services from a single provider, which may create additional technical cost for the customer. For example, even if Google BigQuery is able to load data from (‘talk to’) Azure Storage, more steps may be required to load data from Azure Storage into Google BigQuery compared to loading data from Google’s storage service.²⁹³
- 5.25 We recognise that some of the multi-cloud services offered by the hyperscalers may alleviate this technical effort. However, we have received limited evidence of customers adopting an integrated multi-cloud architecture. Customers have also highlighted technical challenges (as well as financial costs) associated with this type of deployment. In particular:

²⁹² For example, the customer may need to procure and deploy an adaptor to make cloud services from different providers ‘talk to each other’ (i.e. interoperate).

²⁹³ Customers will face additionally administrative burdens related to cross-cloud communication and access, i.e. they will have to enable network connectivity and grant the necessary permissions in both clouds. In many cases, customers will be required to deploy a custom script, application or introduce a new workflow which will act as an intermediary layer between the two services, bridging the interoperability gaps.

- a) As discussed in Annex 7, it is difficult to use the findings of the quantitative part of the market research to assess the prevalence of integrated multi-cloud. 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 somewhat integrated multi-cloud set-up, but these companies tend to be smaller and have simpler tech requirements.²⁹⁴
- 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 them to focus on the areas where their business adds value.²⁹⁵
- c) Some smaller cloud providers mentioned that lack of interoperability between clouds can act as a barrier to integrated multi-cloud and hinder their ability to acquire customers.²⁹⁶

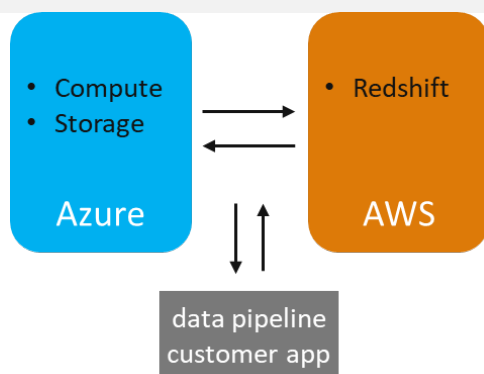
5.26 Overall, we provisionally consider that technical barriers to deploy an integrated multi-cloud architecture are likely to be material. These barriers mostly stem from a lack of interoperability between clouds, which are likely to require material effort from customers integrating services from different cloud providers.

²⁹⁴ Context Consulting research report, slide 80.

²⁹⁵ [redacted] cited limitations to or lack of interoperability as one barrier to integrated multi-cloud ([redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]); [redacted] 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 ([redacted] response dated [redacted] to our customer questionnaire, question [redacted]); [redacted] said that the limitation of interoperability is the main barrier to integrated multi-cloud ([redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]); [redacted] said there is lack of interoperability and/or difficulty in incorporating services from multiple providers in a single cloud architecture especially at the speed ([redacted]); [redacted] cited limitations to or a lack of interoperability as one of the main barriers to integrated multi-cloud ([redacted] response dated [redacted] to the s.174 notice dated [redacted], questions [redacted]); [redacted] said that integrated multicloud requires the use of an additional abstractions layer which risks access to innovation, and leads to higher development costs and worse performance ([redacted] response dated [redacted] to our customer questionnaire, question [redacted]); [redacted] said that the main barrier to integrated multicloud is the effort to configure two different environments ([redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]); [redacted] ranked limitations to or lack of interoperability as the top barrier to integrated multi-cloud ([redacted] response dated [redacted] to our customer questionnaire, question [redacted]); [redacted] said complexity in terms of interoperability has been key factor in it not having adopted integrated multi-cloud ([redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]).

²⁹⁶ 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); [redacted] response dated [redacted] to the s. 174 notice dated [redacted], pages [redacted] and questions [redacted].

Box 5.3: Hypothetical example of customer deploying integrated multi-cloud



A digital advertising agency is running its main operation in Microsoft Azure where they use compute, storage and various other cloud services to deliver customised advertisements to their customers. When processing the resulting data to assess performance, they want to use the Amazon RedShift data warehouse among other analytics and machine-learning services offered by the AWS cloud.

Their first technical challenge is to allow the agency’s operations team to manage a second cloud provider without creating significant overheads. For that, the agency opts to expand their existing cloud-agnostic management tool by writing an adaptor for AWS. Since data cannot be loaded into Amazon RedShift directly from Azure, the agency writes an application which periodically copies data to Amazon S3 and stores it there in a RedShift compatible data format.

(b) Interoperability and portability considerations when duplicating clouds

- 5.27 Cloud duplication 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 work equivalently on all of them.
- 5.28 To implement this multi-cloud architecture, customers need to make sure that their relevant applications can be executed on both clouds. Where there is a low degree of interoperability and portability, deploying a customer application on two distinct clouds requires extensive reconfiguration effort by costumers. The scale of these challenges varies depending on the use-case. However, the main factors that determine the level of technical effort are usually the number of applications that need to run on both clouds and the tightness of their integration into the proprietary cloud services of the primary cloud (i.e. the more proprietary cloud services the customer is using on their primary cloud, the bigger the reconfiguration effort needed to replicate these services on the back-up cloud).²⁹⁷
- 5.29 The market research and the responses to our customer questionnaire indicate that cloud-duplication is relatively infrequent and typically adopted for resiliency reasons and/or to meet regulatory requirements. Customers may only duplicate some parts of their cloud

²⁹⁷ The Government’s guidance on managing technical lock-in in the cloud states that technical lock-in can be caused by too much tight integration with provider-specific services or products (see [Managing technical lock-in in the cloud - GOV.UK \(www.gov.uk\)](https://www.gov.uk/government/guidance/managing-technical-lock-in-in-the-cloud)). This is also confirmed by evidence we have received from stakeholders. For example, [redacted] acknowledged that it would be very difficult for them to switch cloud provider given the large number of PaaS products they have invested in (Ofcom/[redacted] meeting, [redacted]); [redacted] said that external facing workloads would have a significant technical barrier as they utilize more cloud native services ([redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]); Ofcom/[redacted] meeting, [redacted]; Ofcom/[redacted] meeting, [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

architecture (as opposed to the full architecture) to maximise service availability for critical applications in case of outage. In particular:

- a) The market research suggests that a small minority of customers (15% of those who use more than one IaaS/PaaS provider) have implemented a multi-cloud architecture where they have a primary IaaS/PaaS provider and secondary one as a back-up.²⁹⁸ Around 64% of respondents who adopted this failover multi-cloud architecture indicated either a lack of interoperability or other technical challenges as the main difficulty.^{299, 300}
- b) The responses to our customer questionnaire indicate that mirroring cloud architectures with the aim to run the same application across clouds is typically adopted for resiliency reasons and/or to meet regulatory requirements. For example, [X] uses a multi-cloud architecture to duplicate their databases on Google Cloud which would allow the company to maintain critical functionality if anything were to happen to its primary cloud, AWS. It said that this choice has been made from both a strategic and regulatory perspective.³⁰¹

5.30 Our evidence also indicates that customers wishing to duplicate, all or some parts of, their cloud architectures face substantial technical and financial costs associated with setting up and maintaining two parallel clouds. This is likely to further discourage adoption of this specific multi-cloud set-up. For example:

- a) The market research found that 39% of customers adopting cloud duplication for back-up purposes indicated lower cost efficiency as an implementation challenge.³⁰²
- b) Large organisations responding to our customer questionnaire indicated cost was a key barrier. For example, [X] explained that it would be cost prohibitive to maintain both clouds actively all the time. Therefore, any multi-cloud alternative would need to be limited but with the ability to scale up rapidly.³⁰³ [X] indicated replication of services as the most important obstacle for multi-cloud as it drives cost and complexity of the overall environment.³⁰⁴ BT Group said that operating two cloud environments concurrently when switching critical services represents an important cost-related consideration. While this submission was not referring to a cloud duplication case it supports the view that replicating clouds may be particularly costly.³⁰⁵

5.31 We note that the same technical challenges largely apply to ISVs who wish to duplicate their services on different public clouds with the aim to access a larger pool of users.

²⁹⁸ Context Consulting research data tables, Q29. This figure is 14% after excluding respondents using only private cloud.

²⁹⁹ More specifically, 38% of respondents mentioned lack of interoperability between clouds, and 36% of respondents mentioned technological difficulties. Context Consulting research data tables, Q31.

³⁰⁰ The market research also suggests that around 40% of customers using more than one IaaS/PaaS provider (39% if we exclude users of private cloud only) are 'spreading similar workloads across IaaS/PaaS providers'. As explained in Annex 7, it is difficult to assign these responses to a specific multi-cloud architecture.

³⁰¹ [X].

³⁰² Analysis of Context Consulting research data, Q31.

³⁰³ [X].

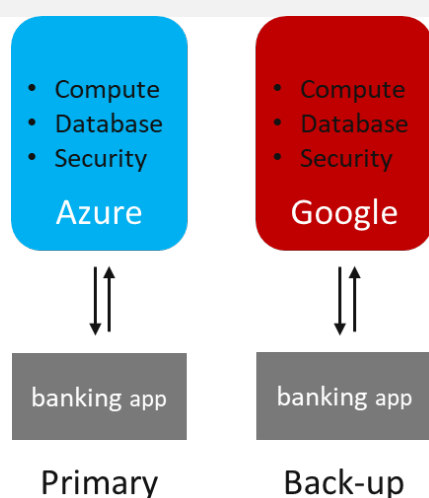
³⁰⁴ [X] response dated [X] to our customer questionnaire, question [X].

³⁰⁵ BT Group response to the CFI, page 9.

Hence, a low degree of interoperability and portability also has the effect of materially increasing ISVs' technical effort to deploy their cloud services on multiple clouds. Evidence from ISVs indicates that integrating with additional cloud providers would require material cost and time, which is likely to reinforce network effects.

- 5.32 Overall, we provisionally consider that, due to a low degree of interoperability and portability, customers and ISVs are likely to face some technical challenges when replicating their architectures on multiple clouds. Such challenges are likely to be particularly substantial for customers who wish to mirror more parts of their architecture across clouds, especially where such an architecture is tightly integrated with proprietary cloud services.

Box 5.4: Hypothetical example of customer duplicating clouds for resilience



A small online payment provider wants to ensure that customers are still able to make payments during an outage. To minimise cost and complexity, they decide to replicate a part of their existing architecture to a different cloud provider to improve reliability and resilience for the most essential services. Instead of mirroring everything, they agree to deploy a simplified product version with reduced functionality. They also aim to operate their spare cloud instances in standby mode for failover purposes. Instead of being constantly online, they will only be activated and enabled during outages, when the main application becomes unavailable.

To facilitate the new application architecture, the payment provider needs to redesign and engineer the relevant services and adopt cloud neutral software principles to allow them to run in different cloud environments. Additionally, the operational complexity of keeping software in multiple clouds requires extra steps for continuous integration and coordination between monitoring systems.

(c) Interoperability and portability considerations when using multiple siloed clouds

- 5.33 Some customers may run different applications, store different sets of data and/or use different cloud services hosted on two or more public clouds with no or minimal integration between these clouds (i.e. different customer applications and data are siloed on different clouds).
- 5.34 As with other multi-cloud architectures, some technical effort is required to implement a siloed multi-cloud architecture. In particular, to the extent that siloed customer applications and data use the same functionalities across clouds (e.g. storage, security, encryption, application management and billing), customers have to duplicate their effort

to set up and manage the cloud services that provide these functionalities.³⁰⁶ This technical effort would be more acute where such services have a low level of standardisation.

- 5.35 However, the technical effort is likely to be more contained compared to other types of multi-cloud architectures. This is because unlike the cloud duplication scenario, not all functionalities would need to be duplicated. Moreover, unlike the integrated multi-cloud scenario, customer applications, customer data and cloud services hosted on different clouds in a siloed way have minimal to no interactions.
- 5.36 The evidence received indicates that, where it occurs, multi-cloud is predominantly siloed as this is less technically demanding and more cost efficient than integrated multi-cloud. In particular:
- a) According to the market research, siloed multi-cloud is the most frequently adopted architecture by customers using multiple IaaS/PaaS providers (45%).³⁰⁷ However, this 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.³⁰⁸ 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.³⁰⁹ This suggests that in some circumstances siloed multi-cloud may be chosen as an easier – and yet unpreferred – multi-cloud architecture, from customers wanting to procure specific services from different cloud providers.
 - b) [X] and [X] submitted that they have adopted a siloed multi-cloud architecture as running workloads across clouds would be inefficient³¹⁰ or impose additional technical or commercial costs.³¹¹ [X] and [X] did not cite any particular technical challenge with adopting a siloed multi-cloud approach.³¹²
- 5.37 In addition, adoption of siloed multi-cloud may also reflect an uncoordinated process of migration to the public clouds within a single organisation. This is where different departments or different subsidiaries may have independently migrated to different cloud providers possibly across different geographical areas. For example, [X] said it has no firm-wide strategy in relation to use of cloud infrastructure services and so internal teams can decide on what approach suits them best. Therefore, whilst applications are mostly siloed within a single cloud, this is largely organised around department considerations,

³⁰⁶ For example, a customer may need to set up security groups, set up firewall rules and integrate with monitoring tools twice.

³⁰⁷ This figure is 47% after excluding respondents using only private cloud (Context Consulting research data tables, Q29). As discussed in Annex 7, 40% of customers (39% excluding users of private cloud only) who ‘spread similar workloads across clouds’ may capture some siloed multi-cloud.

³⁰⁸ 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.

³⁰⁹ Context Consulting research report, slide 80.

³¹⁰ [X] response dated [X] to our customer questionnaire, question [X].

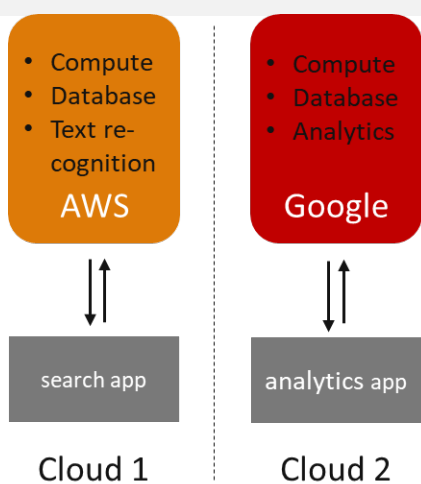
³¹¹ [X] response dated [X] to our customer questionnaire, question [X].

³¹² [X] response dated [X] to the s.174 notice dated [X], question [X]; and [X] response dated [X] to the s.174 notice dated [X], question [X].

such as previous deployments and skills sets, rather than distinct workloads for specific cloud providers.³¹³

- 5.38 Overall, our preliminary view is that technical barriers also exist for customers wanting to adopt a siloed multi-cloud architecture. These barriers mostly stem from duplicating the effort to set-up and manage comparable services across clouds and will be more acute where such services have a high level of technical differentiation. However, our preliminary assessment is that these challenges are likely to be generally lower compared to other multi-cloud architectures. This is because not all functionalities would need to be duplicated and cloud services hosted on different clouds would have minimal to no interactions.
- 5.39 The evidence also suggests that some customers would ideally want to integrate services from different clouds (i.e. implement integrated multi-cloud) but settle for a siloed multi-cloud deployment due to the higher technical challenges associated with integrated multi-cloud. As such, the relatively higher levels of take-up for siloed multi-cloud might in fact signal the existence of potential demand for integrated multi-cloud.

Box 5.5: Hypothetical example of a customer deploying siloed multi-cloud



A large telecommunications vendor has multiple organisational divisions with only little overlap in their activities which requires no integration between their services. Each division has their own dedicated technical resources and significant autonomy when it comes to architectural infrastructure decisions and can therefore act much like an independent entity. Because interoperability is not a concern, each organisational unit can create their self-contained silo within the cloud of their choice.

In practice, the teams will never have complete autonomy and there will be some top-down direction or technical framework that will have to be followed. Equally, there will be some limited integrations between the silos (e.g. to enable centralised authentication or auditing) but those are usually relatively basic and insignificant in terms of complexity and data volumes.

Switching: the evidence suggests that a lack of interoperability and portability can materially increase barriers to switching

- 5.40 In this subsection we consider the extent to which a lack of interoperability and portability may hinder customers' ability to switch:

³¹³ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

- a) **Between public clouds:** customers switch from using cloud A (origin cloud) to using cloud B (target cloud). The switch can be related to all of the cloud services (full switch) or some of the cloud services (partial switch) that the customers use on the origin cloud.
- b) **Within the same public cloud:** customers switch from using a cloud provider’s services to using ISVs’ services hosted on the same public cloud (or vice versa).

- 5.41 The market research found that 18% of customers switched IaaS/PaaS providers.³¹⁴ Some customers may have understood switching to relate to switching between different public clouds, others may have understood it to mean switching within the same public cloud or switching between on-premises IT/private cloud and public cloud. This is consistent with the findings of the qualitative part of the market research which encountered few if any examples of organisations switching away from one of the hyperscalers. Therefore, we consider this figure may overstate the prevalence of switching between clouds and within the same cloud.
- 5.42 We recognise that the presence of multiple types of switching may also have an impact on the market research results around specific barriers to switching. However, we consider this impact to be limited with regards to technical barriers as the nature of these barriers is likely to be broadly similar across all types of switching. As such, we consider the market research results are a reasonable indication of the potential technical issues in switching between public clouds and within clouds, and we have therefore set them out in the relevant section below.
- 5.43 Overall, the evidence received indicates that switching is likely to require material effort, but its scale will depend on the actual switching scenario and use-case. This is summarised in the table below and discussed in more detail in the following paragraphs.

Table 5.6: Overview of preliminary findings on the technical effort to switch

| Switching scenario | Technical barriers | Customers most affected by barriers | Level of effort |
|------------------------------------|--|---|-----------------|
| a) Switching between clouds | <p>Technical effort to reconfigure all applications and data to work on the target cloud.</p> <p>Technical effort to set-up and maintain a multi-cloud architecture if switching some services only (see barriers to multi-cloud).</p> | <p>Customers wishing to port many applications.</p> <p>Customers with architectures tightly integrated with proprietary cloud services.</p> | High |
| b) Switching within cloud | <p>Technical effort to reconfigure applications and data to ensure they work with the target cloud services.</p> | <p>Customers wishing to port many applications.</p> <p>Customers with architectures tightly integrated with proprietary cloud services.</p> | Medium |

Source: Ofcom.

³¹⁴ Context Consulting research data tables, Q47. This figure is 21% after excluding respondents who only use private cloud.

a) Interoperability and portability considerations when switching between public clouds

- 5.44 Switching between clouds involves customers recreating their cloud architecture on 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 one part of it only (partial switch). Hence, the switch can be related to all or some of the cloud services that customers use on the origin cloud.
- 5.45 When switching between public clouds, customers need to ensure their applications and data work and perform equivalent tasks in the target cloud. 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 low degree of interoperability and portability 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 will likely vary depending on the use-case. Similar to cloud duplication, 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.
- 5.46 Compared to a full switch, in a partial switch the scale of the technical challenges is likely to be lower, as the customer is only changing some parts of their cloud architecture. However, this case can effectively be regarded as a switch from a single-cloud to a multi-cloud architecture which, as discussed above, may come with additional interoperability and portability challenges (e.g. effort to ensure that services from different clouds can ‘talk to each other’). Overall, the degree of technical effort to implement a partial switch between clouds will likely be substantial but its exact scale will depend on the number of services being switched, the tightness of their integration into the proprietary services of the origin cloud and the extent to which the origin and the target clouds need to be integrated (i.e. to interoperate).
- 5.47 We acknowledge that customers could employ best practices and tools to reduce barriers to switching (e.g. by building their solutions based on cloud-neutral design principles).³¹⁵ However, the evidence we have received, and our own analysis, indicates that – despite the existence of these mitigation strategies – technical challenges are likely to be material for some customers and use-cases. In particular:
- a) The qualitative part of the market research encountered few, if any, examples of organisations switching away from one of the hyperscalers. However, given that the quantitative part of the market research suggests that a minority of customers (c. 20%) have switched, we consider that switching is possible for some customers. As discussed below, this may be the case for smaller customers with fewer applications and simpler use-cases.
 - b) Among all respondents, application portability was the third most frequently cited barrier to switching (31%), followed by data portability in fourth place (30%) and

³¹⁵ See Annex 7 for further details on this.

interoperability challenges in eighth place (21%).³¹⁶ Overall, 58% of customers indicated at least one of these technical challenges as a barrier to switching. As discussed above, this may understate the overall importance of technical barriers as it excludes ‘time and costs’, which was the most frequently cited barrier to switching (43%)³¹⁷ and may also reflect the time and cost needed to reconfigure data and applications.

- c) 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.³¹⁸ For example, [X] said that while its cloud strategy mandates use of open standards and open-APIs, it has experienced substantial technical barriers to switching as moving supplier means re-architecting applications even if they are abstracted behind IaaS or CaaS (i.e. containers). [X] mentioned it followed a standard migration process, and this has been ongoing for 6 months.³¹⁹
- d) 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.³²⁰

5.48 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

³¹⁶ Context Consulting research data tables, Q52. After excluding respondents using only private cloud, these figures change as follows: application portability (34%); data portability (30%); and interoperability (23%).

³¹⁷ Context Consulting research data tables, Q52. This figure is 43% after excluding respondents using only private cloud.

³¹⁸ [X] said that switching cloud providers will be costly and require extensive rework, implementation and testing ([X] response dated [X] to the s.174 notice dated [X], question [X]); [X] 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 ([X] response dated [X] to our customer questionnaire, question [X]); [X] 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 ([X] response dated [X] to the s.174 notice dated [X], question [X]); 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); [X] 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 ([X] response dated [X] to the s.174 notice dated [X], question [X]); [X] 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 ([X] response dated [X] to our customer questionnaire, question [X]); [X] ranked technical barriers as the second most important barrier to switching ([X] response dated [X] to the s.174 notice dated [X], question [X]); [X] 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 ([X]); [X] said that limitations to or a lack of interoperability is one of the factors preventing it from switching ([X] response dated [X] to the s.174 notice dated [X], question [X]); [X] said that barriers and costs vary but the technical barriers would likely be relatively significant ([X] response dated [X] to the s.174 notice dated [X], question [X]); [X] ranked technical difficulties as the second most important barrier to switching after vendor’s product performance and functionality ([X] response dated [X] to our customer questionnaire, question [X]); [X] expected lack of interoperability to be its biggest problem for migration ([X] response dated [X] to the s.174 notice dated [X], question [X]); [X] said that technical barriers to switching are significant and added only one more major obstacle - data egress fees ([X] response dated [X] to the s.174 notice dated [X], question [X]); 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).

³¹⁹ [X] response dated [X] to the s.174 notice dated [X], questions [X].

³²⁰ BT Group response to the CFI, pages 8-10; [X] response dated [X] to the s.174 notice dated [X], pages [X] and questions [X]; [X] response dated [X] to the s.174 notice dated [X], question [X]; [X] response dated [X] to the s.174 notice dated [X], questions [X].

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 did not find substantial differences in how often customers from different sectors cite technical difficulties as a barrier to switching (around a third each).³²¹
- b) However, the market research suggests that the importance of barriers to switching may vary more by company maturity. For example, 49% of the companies established for less than 2 years indicated at least one 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.³²² This may be because more established companies have more cloud applications.
- c) Similarly, barriers to switching may vary by companies' 'life stage'. For example, 54% of newer companies (i.e., those who identified as start-ups and as companies in post-start-up growth stage) indicated at least one of the following technical challenges as a barrier to switching: data portability, application portability and interoperability. This figure increases to an average of 59% for more mature companies (i.e. companies that didn't identify as 'newer'). The difference is even starker after excluding users of private cloud only, with 49% of newer companies indicating at least one of those technical barriers compared to an average of 61% for more mature companies.³²³ This may be because more mature customers have more cloud applications.
- d) The market research also found some differences in how customers of different sizes perceive technical barriers to switch. 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).³²⁴ We also found that larger

³²¹ The sectors covered by the market research are: construction; finance & insurance; healthcare; IT & technology; logistics, transportation & distribution; manufacturing; professional services; public sector; and retail and wholesale.

³²² Analysis of Context Consulting research data, 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 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.

³²³ Analysis of Context Consulting research data, Q52. More specifically, the percentage of customers indicating at least one of these technical challenges are as follows: 54% of newer companies (49% after excluding users of private cloud only), 54% of businesses currently stabilising (54% after excluding users of private cloud only), 61% of businesses currently investing in expansion (66% after excluding users of private cloud only), 59% of mature businesses currently ticking over (63% after excluding users of private cloud only).

³²⁴ Analysis of Context Consulting research data, 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).

organisations tend to have more cloud use cases.³²⁵ Taken together these findings may 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.³²⁶

- e) Relatedly, the market research indicates that customers with more use-cases were generally more likely to indicate technical challenges as barriers to switching, which further corroborates the view that the degree of technical effort may positively correlate with the number of applications to port. On average 47% of customers with up to three use-cases selected at least one of these technical challenges as a barrier to switching: data portability, application portability and interoperability. This percentage raises to 60% for customers with 4-10 use-cases.³²⁷ Moreover, the market research found that the vast majority of customers (84%) have between 4 and 10 use cases and may therefore face these technical challenges.³²⁸
- f) 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.³²⁹ We understand that, for these customers, cloud-neutral design often becomes simply impractical, due to the additional complexity, general time constraints and the lack of centralised coordination.

5.49 Overall, we provisionally consider that customers are always likely to face some technical barriers when switching between clouds. These barriers mostly stem from a lack of interoperability and portability which tend to increase the reconfiguration effort by customers when porting applications or components to a different cloud. Such effort is 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.

³²⁵ Context Consulting research report, slide 37.

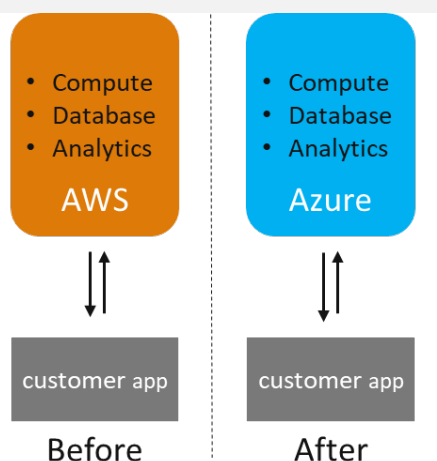
³²⁶ For example, customers with more staff also have more complex identity and access management.

³²⁷ Analysis of Context Consulting research data, Q52. 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.

³²⁸ Context Consulting research report, slide 37.

³²⁹ See paragraph 5.47c).

Box 5.7: Hypothetical examples of customers switching public clouds



Switching the cloud provider of a single **blog application** which relies on a commonly used software package or stack may be an easy endeavour and may be completed within a day or two by qualified personnel.

Similarly, a **small start-up** which operates everything in containers and stringently applies cloud-neutral software architecture and engineering practices may only experience a few minor hurdles when attempting to switch their provider. Depending on service level guarantees of their applications, such

migrations may still involve careful planning stages and may require diligent roll-out preparation, beyond the pure engineering effort. Nevertheless, companies at this scale that operate only a few applications, may be able to accomplish their switch within weeks.

The challenges are more significant when a **customer operates a large number and variety of applications**, often across multiple teams. Even for small businesses with only a few dozen staff, cloud-neutral design is often impractical, due to the additional complexity, general time constraints and the lack of centralised coordination. Over time, such customers become highly embedded into the cloud ecosystem. A customer who has built an interconnected chain of applications, data and analytic pipelines using proprietary cloud functions, database, warehouse, storage, monitoring and integration services (e.g. queues or API gateways), will face significant technical challenges to port those to a different cloud provider. Additionally, when switching cloud providers, data and applications often need to be ported all at once, which significantly increases the scale of the challenge. A migration of that scale may take several months to plan and then several more months and often years to execute. During this time, the customer's technical and operational capacity to service its users would be either severely limited or come to a complete halt, as every resource will be affected by the migration.

Interoperability and portability considerations when switching within the same public cloud

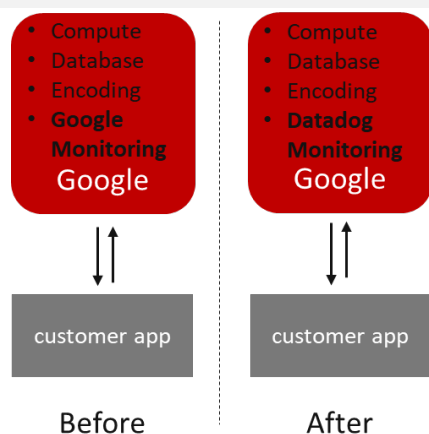
- 5.50 When running applications or storing data on a given cloud, customers can typically choose whether to rely on a cloud provider's or ISV's services on the hosting cloud. If a customer decides to switch between these cloud services, they will need to ensure their applications and data can continue to work as well and perform equivalent tasks.
- 5.51 Similar to the previous scenario (i.e. switching between public clouds), customers switching within the same public cloud would likely need to carry out some reconfiguration work, which materially increases where there is a low degree of interoperability and portability. However, the scale of such reconfiguration effort is likely to be lower compared to the switching between clouds as the customer would only be changing individual cloud infrastructure services while staying within the same cloud.

- 5.52 For example, a customer using AWS that is switching from AWS CloudWatch to the ISV product Datadog Enterprise (all hosted on AWS) would need to make various operational changes and potentially re-engineer some parts of their applications so that they can ‘talk to’ the target cloud service and perform the same tasks. However, these changes are likely to affect fewer parts of their code (e.g. those responsible for collecting metrics as well as connecting, authenticating and exchanging data with the target service). Moreover, since the target service is hosted on the same cloud, many parts of the customers’ applications will likely be unaffected by the switching.
- 5.53 In line with the above, the evidence we have received, indicates that barriers to this type of switching may be lower compared to a full or partial switch between clouds:
- a) One ISV ([redacted]) explained that the only technical requirement for a customer to switch to its services is data migration (i.e. porting data). Data migration can be complex, but it is less so when the data is migrated from one of the cloud providers rather than from the customer’s on-premises solution.³³⁰ Nonetheless, as outlined below, some ISVs ([redacted]) said that switching between a proprietary service to an equivalent service offered by a third party would still require customers to rewrite much of their code.
 - b) We have also received some feedback from customers that is consistent with the above. Specifically, an AWS customer ([redacted]) said that, while it has not switched public clouds, it has switched from third-party database and container services to equivalent AWS services, suggesting that barriers to switch may be relatively lower in this case.³³¹
- 5.54 Overall, we provisionally consider that technical barriers to switching within a cloud may be less acute compared to the barriers to switching between clouds. This is because in this case customers would only be changing some cloud infrastructure services while staying within the same cloud. However, customers are still likely to face some challenges due to the reconfiguration effort needed when switching between poorly interoperable services hosted on the same cloud.

³³⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

³³¹ Ofcom/[redacted] meeting, [redacted]; and Ofcom/[redacted] meeting, [redacted].

Box 5.8: Hypothetical example of a customer switching within a public cloud



In this scenario an online communications network is looking to switch their monitoring system from a first-party solution to a third-party product, provided by an ISV.

First, they set up the target monitoring system, perform the obligatory administrative tasks, establish network connectivity, identity management, and conduct other activities to mirror the existing settings and allow for a smooth transition. Then, they integrate the ISV’s monitoring framework into the applications and prepare new

versions for release. Since they require some legacy data from the old monitoring service, they design and run a process to port and carry over the data to the new system. Finally, they release the new application versions to make the final switch.

Developing cloud-specific skills can add to the technical barriers to switching and multi-cloud

- 5.55 Different cloud providers use different proprietary cloud technologies (e.g. APIs, protocols, workflows, programming languages and data formats). This requires customers to develop specific skills for each cloud they would like to use, creating a risk of ‘skills lock-in’ as a barrier to switching and multi-cloud.
- 5.56 Skills lock-in refers to the situation in which a customer’s cloud technical staff have developed specialised skills applicable to a particular provider’s cloud environment such that it becomes difficult for them to multi-cloud or switch to a different provider. For example, cloud developers or engineers may have become proficient in using a particular provider’s workflows, tools and APIs. If the company 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.
- 5.57 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.³³² One respondent in 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.³³³ [3] provided anecdotal evidence that for the most part people tend to specialise towards one cloud provider and may not be interested in retraining for another.³³⁴

³³² Microsoft website. [Microsoft Certified: Azure Fundamentals](#) [accessed 1 March 2023]; AWS website. [AWS Certification](#) [accessed 1 March 2023]; Google website. [Grow Skills with Google Cloud Training](#) [accessed 1 March 2023].

³³³ Context Consulting research report, slide 118.

³³⁴ [3] response dated [3] to the s.174 notice dated [3], question [3].

- 5.58 Companies looking to switch to or add a new 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 is likely to be more problematic for customers in the public sector or healthcare industry (45% and 43% respectively said that the need to retrain staff is a barrier to switching). 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 indicate that the need to retrain or hire new staff can add to the effort required to switch or add a cloud provider.³³⁵
- 5.59 Skills challenges can be exacerbated by periods of labour market tightness in the tech sector, with customers competing against global tech giants (including the hyperscalers) for talent. One customer ([redacted]) 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 skilled software engineers.³³⁶ 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.³³⁷ We recognise that there are currently a lot of layoffs in the tech sector, which may alleviate the hiring challenges caused by skills shortages.³³⁸ 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.
- 5.60 We note that the hyperscalers provide some free training and are making efforts to improve their training offerings, [redacted],³³⁹ 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.³⁴⁰ [redacted].³⁴¹ However, these training schemes are likely to be more valuable to customers new to the cloud, rather than those looking to switch or add a new provider, who still face opportunity costs in terms of time and effort to retrain their workforce.

³³⁵ See Annex 7.

³³⁶ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; also see [Forbes, June 2021](#) [accessed 14 February 2023].

³³⁷ Analysys Mason, 2023. [Data center network automation for CSPs: key trends, challenges and requirements for the cloud-native networking era](#) [accessed 14 February 2023]

³³⁸ See for example Computer World, February 2023. [How layoffs at Google could affect enterprise cloud services](#) [accessed 14 February 2023].

³³⁹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

³⁴⁰ Microsoft blog, 2020. [How Microsoft helps customers adopt Azure through developer education](#) [accessed 1 March 2023]

³⁴¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

We have seen evidence that technical barriers may be exacerbated by specific business practices

5.61 The evidence we have received suggests that some of the barriers to switching and to the adoption of multi-cloud may be a result of specific practices of the hyperscalers – and particularly of AWS and Microsoft – which tend to reduce interoperability and portability. Based on the evidence we have received, such practices can be grouped into two broad categories:

- a) use of proprietary cloud technologies when building upon open-source software and open standards; and
- b) restrictions to cross-cloud interoperability.

5.62 In the following paragraphs we consider the extent to which these practices may be driving some of the technical effort required to switch and multi-cloud. We acknowledge that some of these practices may not have been designed with the intention of increasing barriers to switching and multi-cloud and may be beneficial to customers. For example, cloud providers may design their systems in this way because it allows them to introduce certain innovations (e.g. offer certain new features).

Use of proprietary cloud technologies when building upon open-source software and open standards

5.63 We have received submissions suggesting that all hyperscalers, but particularly AWS and Microsoft, may be limiting the interoperability of their cloud services by using proprietary cloud technologies (e.g. proprietary APIs). According to these submissions, in some cases these hyperscalers might initially build their cloud services on open cloud technologies (e.g. open standards and open APIs) but then tweak them (for example, changing some parts of the APIs, or changing certain features or functionalities). As a result customers using these services may need to rewrite some of their code if they wish to switch or multi-cloud.³⁴² More specifically:

- a) An ISV ([redacted]) explained that AWS and Microsoft have developed proprietary [redacted] services which they sell as compatible to the [redacted] 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 the hyperscalers – that these services are fully compatible with (or even a version of) the [redacted] service. However, according to [redacted], the compatibility of these first-party services is actually quite limited, meaning that a customer using such [redacted] compatible services would need to considerably reconfigure their application if they wanted to switch to [redacted].³⁴³
- b) Another ISV ([redacted]) explained that the hyperscalers build their services on open standards and open APIs but tweak them so as to make it more difficult for customers to switch as they would need to rewrite much of their code. They explained that this

³⁴² As noted later in this section, while the hyperscalers may also offer the original version of the technology, customers may be induced to take-up these first-party cloud services (instead of the original open-source) 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.

³⁴³ Ofcom/[redacted] meeting, [redacted].

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.³⁴⁴

- c) A cloud provider ([redacted]) 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. [redacted] 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.³⁴⁵
- d) OVHcloud submitted that any customer willing to switch from a hyperscaler to an alternative cloud provider faces strong hurdles as a result of “lock-in” strategies implemented by those operators. These include the hyperscalers’ use of differentiated features, functionalities and APIs which may slow the migration towards an alternative cloud provider as it would force customers to reconfigure some parts of their code.³⁴⁶
- e) IBM submitted that interoperability and portability are critical to avoid vendor lock-in and reduce switching costs between the different cloud platforms. Use of open and widely used standard can improve interoperability and portability. However, substantial differences between the cloud offerings of different providers persist, limiting interoperability and thus competition.³⁴⁷

Restrictions to cross-cloud interoperability

- 5.64 We also have seen evidence that some of AWS and Microsoft cloud services can only be used within their respective public clouds. When this is the case, customers are prevented from building or switching to an integrated multi-cloud architecture involving such services. For example, if an analytics service offered by cloud A can only load data from a storage service offered by the same cloud, customers will not be able to set-up or switch to³⁴⁸ an integrated multi-cloud architecture combining the analytics service offered by cloud A with the storage service offered by cloud B.³⁴⁹
- 5.65 Customers’ inability to switch and multi-cloud may in turn reinforce barriers to entry and expansion for new or smaller providers who may find it more challenging to attract customers to their clouds. For example, absent these restrictions, a smaller provider may be able to persuade customers to use their storage service in combination with the

³⁴⁴ [redacted]; and Ofcom/ [redacted] meeting, [redacted].

³⁴⁵ [redacted].

³⁴⁶ OVHcloud response dated 17 November 2022 to s.174 notice dated 27 October 2022, questions 15 and 27.

³⁴⁷ IBM response dated 6 December 2022 to s.174 notice dated 25 October 2022, part A question 17, paragraphs 3 and 4, page 16.

³⁴⁸ This would include customers taking up storage and analytics from cloud A who want to switch to the storage solution of a different provider. But also, customers taking up storage and analytics from cloud B who want to switch to the analytics solution of cloud A.

³⁴⁹ 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.

analytics service of a larger provider. With the restrictions in place this will be more difficult as the customer would have to switch their entire cloud architecture which, as discussed above, is likely to require a material effort.

5.66 Consistent with our analysis, the evidence received suggests that these restrictions may hinder customers’ ability to switch and multi-cloud and, as a consequence, compound the effect of barriers to entry and expansion:

- a) A cloud provider ([redacted]) submitted that some cloud providers exclusively provide cloud-related services within their own cloud infrastructure. This mechanism tends to limit the possibilities for a user to migrate to another cloud-provider. The more services from a single provider a customer uses, the more data it generates within that provider’s infrastructure. In turn the more data gathered by the cloud provider, the more complex, costly and time-consuming it becomes to switch to an alternative supplier.³⁵⁰
- b) Another cloud provider ([redacted]) discussed an example of an AWS service (SageMaker) which is only available on AWS cloud. In its view there is a lack of transparency for customers when they first start adopting AWS SageMaker that in order to realise the benefits, it is necessary to consume the wider AWS ecosystem of services, leading to lock-in, a lack of interoperability and preventing users selecting best-of-breed technologies across the market.³⁵¹
- 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. As a result of these limitations, users cannot have different services of different providers work with each other, which hinders the use of multi-cloud.³⁵²

5.67 Based on an analysis of publicly available documentation, we have so far 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 5.9: Examples of AWS and Microsoft cloud services with potential interoperability limits

| Supplier | Service | Customer uptake | Potential limits to interoperability ³⁵³ |
|----------|---------------|----------------------------|---|
| AWS | Amazon Athena | [redacted]. ³⁵⁴ | Can only query data stored on Amazon S3. |

³⁵⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

³⁵¹ Ofcom/ [redacted] meeting, [redacted].

³⁵² The Dutch competition authority (ACM), May 2022, [Market Study Cloud Services](#), page 61.

³⁵³ For AWS see: [Amazon Athena](#); [Amazon Interactive Video Service](#); [Amazon Kinensis Video Streams](#); [Amazon Omics](#); [Amazon Pinpoint](#); [Amazon RedShift](#); [Amazon SageMaker](#); [Amazon SageMaker DataWrangler](#); [Amazon Timestream](#); and [AWS IoT Events Documentation](#). For Microsoft see: [Azure Stream Analytics](#); and [Azure IoT Hub](#).

³⁵⁴ AWS response dated 28 November 2022 to the s.174 notice dated 24 October 2022, part A question 33, Annex A Q33.1.11, p.84.

| Supplier | Service | Customer uptake | Potential limits to interoperability ³⁵³ |
|----------|--------------------------------|--|--|
| AWS | Amazon IVS | This service is used by large companies such as GoPro. ³⁵⁵ | Can only auto-record to Amazon S3. |
| AWS | Amazon Kinesis Video Streams | [redacted]. ³⁵⁶ | Can only deliver extracted images to Amazon S3. |
| AWS | Amazon Omics | AWS website shows this service is popular with customers in the clinical space. ³⁵⁷ | Uses Amazon S3 for data import and export. |
| AWS | Amazon Pinpoint | [redacted]. ³⁵⁸ | Allows adding Amazon Personalize recommendations to a marketing email campaigns, but not from third-party recommendations engines. |
| AWS | Amazon RedShift | \$(redacted) ((redacted)% YoY growth) ³⁵⁹ [redacted]. ³⁶⁰ | Can only bulk load data from Amazon S3. |
| AWS | Amazon SageMaker | \$(redacted) ((redacted)% YoY growth) ³⁶¹ | 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 | \$(redacted) ((redacted)% YoY growth) ³⁶² | Can only use AWS Backup service to manage backups. |
| AWS | AWS IoT Events | \$(redacted) ((redacted)% YoY growth) ³⁶³ | Can only trigger actions with other AWS services. |

³⁵⁵ [Interactive Live Streams – Amazon Interactive Video Service – Amazon Web Services](#)

³⁵⁶ AWS response dated 28 November 2022 to the s.174 notice dated 24 October 2022, part A question 33, Annex A Q33.1.11, p.84.

³⁵⁷ [Genomic Data Analysis – Amazon Omics Customers – Amazon Web Services](#)

³⁵⁸ 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.

³⁵⁹ [redacted]. 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).

³⁶⁰ AWS response dated 28 November 2022 to the s.174 notice dated 24 October 2022, part A question 33 (Annex A Q33.1.11, p.84).

³⁶¹ Ofcom analysis of AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, part B question 4, Annex 2.

³⁶² AWS response dated 28 November 2022 to the s.174 notice dated 24 October 2022, part A question 33, Annex A Q33.1.26, p.38.

³⁶³ Ofcom analysis of AWS response dated 12 December 2022 to the s.174 notice dated 24 October 2022, part B question 4, Annex 2. [redacted].

| Supplier | Service | Customer uptake | Potential limits to interoperability ³⁵³ |
|-----------|------------------------|---------------------------------------|---|
| Microsoft | Azure Stream Analytics | \$(%) ([%] YoY growth) ³⁶⁴ | Exclusively support native Azure services as inputs. |
| Microsoft | IoT Hub | \$(%) ([%] YoY growth) ³⁶⁵ | Allows basic interoperability but prioritises integration with other Azure service. |

Source: Ofcom analysis of AWS and Azure data, internal documents and websites.

- 5.68 AWS and Microsoft submitted that they do not restrict access to their cloud services to third parties, but clarified that:
- a) AWS noted that there are AWS services, or aspects of AWS services, that are only available on AWS because the underlying technology is a unique AWS innovation. For example, AWS offers EC2 instances using AWS Graviton processors that were designed by AWS based on the Arm architecture. While compute instances using Arm architecture are available from other providers (e.g. Oracle Ampere A1 compute instances), workloads running on Graviton processors are available only on AWS.³⁶⁶
 - b) Microsoft suggested that the situation for SaaS offerings can be different. Microsoft said that, like any other company offering cloud powered SaaS solutions, it must make a decision as to whether to rely upon its public cloud or instead to create its own unique full stack SaaS solution. By way of example, Microsoft said it has launched a cloud game streaming service that does not run on Azure infrastructure. Microsoft said Azure does not provide the best platform for such a service and instead it created an end-to-end experience in certain Microsoft data centres using Xbox console hardware. We consider this argument may also apply to PaaS offerings.³⁶⁷
- 5.69 We understand that, in practice, AWS and Microsoft impose such restrictions by not sharing information (e.g. APIs and protocols) required to interoperate with their services. We note that other cloud providers, including Google, offer equivalent services without imposing any such restrictions. This indicates that it is possible for these services to be made technically interoperable with those of other cloud providers. In particular:
- a) Google allows equivalent services to interoperate with other cloud providers. For example, BigQuery, which is comparable to AWS Redshift, can bulk load data from any third-party storage service.
 - b) ISVs generally offer equivalent services on a standalone basis and deploy them across multiple cloud providers. For example, Snowflake – which is comparable to AWS Redshift – can interoperate with storage services of AWS, Google and Microsoft.

³⁶⁴ 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.

³⁶⁵ 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.

³⁶⁶ AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, part A question 31.

³⁶⁷ Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, part A question 31.

5.70 The above evidence suggests that incentives in setting the preferred level of interoperability differ across cloud providers. Challengers and smaller cloud providers may prefer a higher level of interoperability to increase demand by making it easier for customers to switch away from rival clouds and by capturing additional demand from customers wishing to multi-cloud. Cloud providers with an established market position and large existing customer base may prefer a lower level of interoperability to avoid losing customers. In this sense prevailing levels of interoperability can provide an indication of the relative strength of different cloud providers.

Provisional conclusion

- 5.71 We provisionally find that a lack of interoperability and portability can amplify the degree of technical differentiation between clouds and therefore increase technical barriers to switching and multi-cloud. These barriers include: (i) the effort required to reconfigure data and application such that these can work on the target cloud(s); and (ii) the time and costs needed to develop cloud-specific skills.
- 5.72 We consider that technical barriers are likely to be particularly strong for customers wishing to integrate services from multiple clouds (i.e. integrated multi-cloud) or replicate their cloud architecture on a different / additional cloud (i.e. switching or duplicating clouds).
- 5.73 Some customers may be relatively unaffected by these costs. This includes those who use a few applications that have been built on common or sufficiently similar standards. However, our evidence suggests that a significant number of users are concerned about the technical hurdles they face to switch clouds, and some report having settled for a sub-optimal approach to multi-cloud.
- 5.74 Our evidence indicates that the reconfiguration efforts are likely a material barrier to switching and multi-cloud for customers with large numbers of applications and / or with a cloud architecture that is tightly integrated with many first-party proprietary services. Based on the results of the market research, we believe this is likely to encompass a material portion of customers across several sectors and reflect a significant fraction of revenue.
- 5.75 Each of the hyperscalers offer services and tools which can mitigate technical barriers to switching and multi-cloud in principle, including the use of open-source software and open standards which are available across clouds. However, our evidence indicates that uptake of some of these tools is limited.
- 5.76 Moreover, it appears that technical barriers may be exacerbated by practices employed by the hyperscalers and, in particular, by Microsoft and AWS. These include:
- building first-party services on open standards and open APIs but make some adjustments so that the resulting first-party service is ultimately less compatible with the underlying standards; and

- the use of restrictions to interoperability for some AWS and Microsoft cloud services, which may prevent customers from building an integrated multi-cloud architecture involving such services, or switch part of this service to a rival cloud provider.

5.77 Overall, these technical barriers to switching can lessen competition between cloud providers for some existing customers, by lowering the threat of these customers switching all or some of their workloads to benefit from better prices or higher quality cloud services. This could particularly inhibit competitive constraints faced by AWS and Microsoft, whose products may be less interoperable than those of other cloud providers in the market.

Data egress fees

5.78 Customers may need to transfer their data:

- a) into a cloud provider's infrastructure;
- b) within that cloud provider's infrastructure; and
- c) out of that cloud provider's infrastructure.

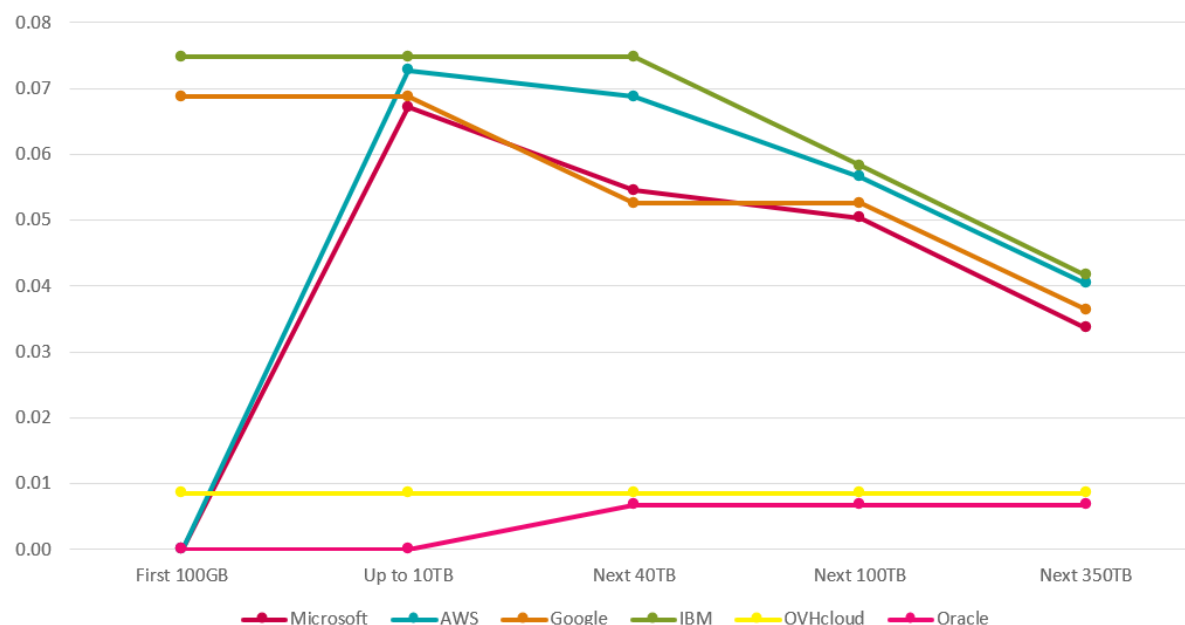
5.79 We understand that transferring data into the cloud, known as ingress, is free.³⁶⁸ However, customers face charges when moving data within their cloud provider's infrastructure, and when they transfer data out of a cloud provider's infrastructure. We refer to charges for transferring data out of a cloud provider's infrastructure as egress fees. These egress fees apply when a customer is transferring data to an end user, when a customer is seeking to use integrated cloud solutions which require data transfers (i.e. to use an integrated multi-cloud architecture), or when a customer is looking to switch from services on one cloud provider's infrastructure to another.

5.80 Figure 5.10 below shows the egress fees charged by six cloud providers for transferring data out of the cloud via the public internet. The figure illustrates how AWS, Google, Microsoft and IBM charge fees which are 5-10 times higher than OVHcloud and Oracle. For egressing via a private network, Google and Microsoft charge higher egress fees than those shown in the figure.³⁶⁹

³⁶⁸ For example, none of the hyperscalers charge for ingress. See Microsoft website, [Pricing – Bandwidth | Microsoft Azure](#) [accessed 8 March 2023]; AWS website [Amazon S3 Simple Storage Service Pricing - Amazon Web Services](#) and [EC2 On-Demand Instance Pricing – Amazon Web Services](#) [accessed 8 March 2023]; Google website. [All networking pricing | Virtual Private Cloud | Google Cloud](#) [accessed 8 March 2023].

³⁶⁹ Both Google (premium tier) and Microsoft (Premium Global Network) charge higher prices for egressing via a private network, which is the default routing they offer to their customers. Google's premium tier egress fees are higher than every other cloud provider shown in the figure except for data transfers between 10-50TB per month. The Microsoft Premium Global Network internet egress fees are almost identical to the AWS public internet egress fees. See [Routing preference in Azure | Microsoft Learn](#) and [All networking pricing | Virtual Private Cloud | Google Cloud](#).

Figure 5.10: Egress fees for transferring data out via the public internet, March 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 8 March 2023]. OVHcloud response dated 17 November 2022 to s.174 notice of 27 October 2022, Q9, p.13.

Notes: 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 a single region.

5.81 Table 5.11 below shows the volume of data transferred out to the public internet from each of AWS and Microsoft Azure in 2021, and the associated revenues received for transferring that data out to the public internet. Egress fees account for [redacted]. Egress volumes are [redacted].³⁷⁰

Table 5.11: Data transferred out to the public internet and associated egress fees received for those data transfers, UK, AWS and Microsoft Azure, 2021

| | AWS | Microsoft Azure |
|--|---------------|--------------------------|
| Data transferred to public internet from UK region | [redacted] GB | [redacted] [300-400m] GB |
| Egress fees received for data transfers out to public internet | £[redacted] | £[redacted] [20-30m] |

³⁷⁰ AWS response dated 13 January 2023 to follow-up email dated 16 December 2022 concerning the s.174 notice dated 24 October 2022, Table 1; and Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, Annex B7.

Sources: AWS response dated 13 January 2023 to follow-up email dated 16 December 2022 concerning the s.174 notice dated 24 October 2022, Table 1, AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, question 6, Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, Annexes B6 and B7.

Customers are concerned by egress fees, especially in relation to multi-cloud

- 5.82 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 both those which are currently using a single cloud provider, and those which are using multiple cloud providers. The level of concern is higher for those who have switched, or considered switching provider, than for those who have not.³⁷¹
- 5.83 Egress fees are likely to be a particularly strong barrier to multi-cloud. The market research 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.³⁷²
- 5.84 Egress fees also act as a barrier to switching between cloud providers. 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.³⁷³

Many customers are seeking alternative cloud architectures to minimise or avoid egress fees

- 5.85 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.³⁷⁴ Egress fees may 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 approach. 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).
- 5.86 We have heard examples of customers using a single-cloud architecture, rather than purchasing service from multiple cloud providers, because of egress fees. For example, [redacted] stated that avoiding egress charges were a main driver of its decision to purchase multiple cloud services from the same provider.³⁷⁵ We are also aware of specific examples

³⁷¹ Context Consulting research data tables, Q63; Context Consulting research report, slide 131.

³⁷² Context Consulting research data tables, Q31; Context Consulting research report, slide 79.

³⁷³ 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, Q52 and Q53.

³⁷⁴ This includes broadcasters and telecoms operators such as ([redacted], BT and Vodafone) and other customers such as ([redacted]). See BT response dated 27 January 2023 to the s.174 notice dated 1 December 2022, Q18, Q20; [redacted] response dated [redacted] to the s.174 notice dated [redacted], questions [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted] question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], questions [redacted]; Vodafone response to the CFI, page 3, questions 4.2 and 4.7; [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; and [redacted] response dated [redacted] to our customer questionnaire, question [redacted].

³⁷⁵ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

of customers moving 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 5.12 below.

Box 5.12: Examples of customer decisions to move workloads from multi-cloud to single-cloud solutions being influenced by egress fees

[X] 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.³⁷⁶

[X] explains that data egress is the most pertinent obstacle to multi-cloud, considering that data egress has influenced its architectural choice [X]. It has moved its [X] to align with its choice of [X], in part to avoid excessive egress fees.³⁷⁷

- 5.87 Even where customers are choosing to use multiple clouds, many are doing so by siloing workloads on different clouds, rather than using an integrated multi-cloud architecture.³⁷⁸ In some cases this is motivated in part by a desire to avoid paying egress fees.³⁷⁹ Other customers have told us they apply elaborate workarounds (which themselves incur costs) to reduce costs incurred through egress fees.³⁸⁰
- 5.88 Certain customers may be required to use multiple cloud providers for resilience purposes, such as financial institutions. In such cases, we understand customers may limit their usage of multiple clouds to the minimum necessary for resilience purposes. For example, using a secondary provider as failover that would continue to provide some minimum level of functionality in the event of an outage.³⁸¹

The strength of egress fees as a barrier to multi-cloud will not be the same for all customers

- 5.89 Egress fees are likely to be a greater barrier to integrated multi-cloud solutions where a significant amount of data needs to be transferred between the services hosted on different clouds. In these cases, customers would incur high egress fees if they were to

³⁷⁶ [X] response dated [X] to the s.174 notice dated [X], page [X].

³⁷⁷ [X] response dated [X] to the s.174 notice dated [X], questions [X], pages [X].

³⁷⁸ 45% of those adopting multi-cloud solutions use different providers for different workloads, and 15% use one as a principal supplier and use additional supplier as backup/failover. Context Consulting research data tables, Q29 and Context Consulting research report, slide 76. Also see Annex 7, paragraph A7.11.

³⁷⁹ [X] noted that individual workloads are located within a single cloud in part to avoid egress fees. See [X] response dated [X] to our customer questionnaire, question [X].

³⁸⁰ For example, some customers are using private networking to try and minimise egress fees. Monzo introduced a private link between AWS and GCP to route its traffic over to avoid NAT gateway fees. It uses a third party to create a virtual router and two connections, one to each cloud network, and it can bridge between data centres they choose. It is also needed to connect the private link into its network, tell traffic for these Ips to go via the Transit gateway and send traffic which hits the transit gateway down the Direct Connect. [Reducing our NAT Gateway cost with private networking between AWS and GCP \(monzo.com\)](#)

³⁸¹ Flexera found that 42% of customers which were using multi-cloud architectures used them for DR/Failover between clouds (Flexera, 2023. 2023 State of the Cloud report, page 20). [X] has developed a back-up facility on one cloud provider [X], but only sufficient to enable it to quickly scale up if its main cloud provider [X] fails. Ofcom/[X] meeting, [X].

- purchase these services from different cloud providers and transfer the data between them.
- 5.90 Any element of data processing by its very nature requires large volumes of data to be copied across the network, and so scenarios where data processing is carried out in a different cloud to that where the data is stored are likely to incur high egress fees. 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.
- 5.91 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.
- 5.92 Customers most likely to be in these scenarios include (but are not limited to) those involved in personal data processing and analytics (e.g. online advertising, financial services, prediction and forecasting, social networks) as well as media businesses dealing with large volumes of media assets, content distribution and delivery. This could also include public sector and healthcare organisations.³⁸² The market research found that 76% of IaaS/PaaS users purchase databases, 50% of IaaS/PaaS users do so for business intelligence and/or big data, and 42% host CPU intensive applications.³⁸³ We also note that data management and analytics services appear to generate the most PaaS revenues for the hyperscalers in aggregate, implying that they are taken by a significant number of customers.³⁸⁴
- 5.93 On the other hand, some workflows are unlikely to require large amounts of data to flow between cloud services, such as when moving text files. Egress fees will be less of a barrier to multi-cloud for these workflows.
- 5.94 We understand that data volumes do not scale linearly with company size. Companies of all sizes may store large volumes of data in the cloud, or undertake workflows which require large data movements between different cloud services. As such, even smaller companies may find that they face potentially high egress fees for their use cases.
- 5.95 The extent to which egress fees acts as a barrier to multi-cloud is likely to depend on how much a customer could save by single-homing, how that relates to the total cloud spend of that customer and the ability for a cloud provider to offset any material egress fees to make it viable for a customer to host integrated cloud services across several clouds.

³⁸² Context Consulting research data tables, Q15 and Context Consulting research report, slide 80.

³⁸³ Context Consulting research report, slide 35 and Section 3, Figure 3.3.

³⁸⁴ See Annex A5, paragraph A5.28.

5.96 We have received limited data to date to explore this question. To illustrate the potential impact that egress fees can have as a barrier to multi-cloud, we use some stylised examples of the data transfers associated with certain workflows, and set out the egress fees which would be incurred if these workflows utilise services hosted on different cloud infrastructures. These are set out in Box 5.13. We show how egress fees grow with the amount of data transferred per month in Figure 5.14. We then set out that egress fees can be significant relative to:

- a) A customer's spending levels for an example use case if all services are hosted on a single cloud – meaning a significant saving for that use case could be achieved by hosting all services associated with that use case on a single cloud; and
- b) A range of customer spending levels that are representative of a large fraction of cloud infrastructure customers.

5.97 We intend to explore this issue further ahead of our final report.

Box 5.13: Egress fees for multi-cloud scenarios

We have considered some illustrative scenarios of different workflows being used in a multi-cloud context to demonstrate the egress fees which customers could incur.

Scenario 1: Low data transfer - moving reports

A company is hosted on one cloud, and uses a machine learning model to generate reports, which it then wants to load into a different cloud. On average, it generates 1000 reports per day, and each report is 10MB. The company therefore transfers 10GB of data per day, and around 300GB per month. This would only incur a relatively small amount of egress fees (less than £25) each month.³⁸⁵

Scenario 2: Medium data transfer – moving data from storage to database

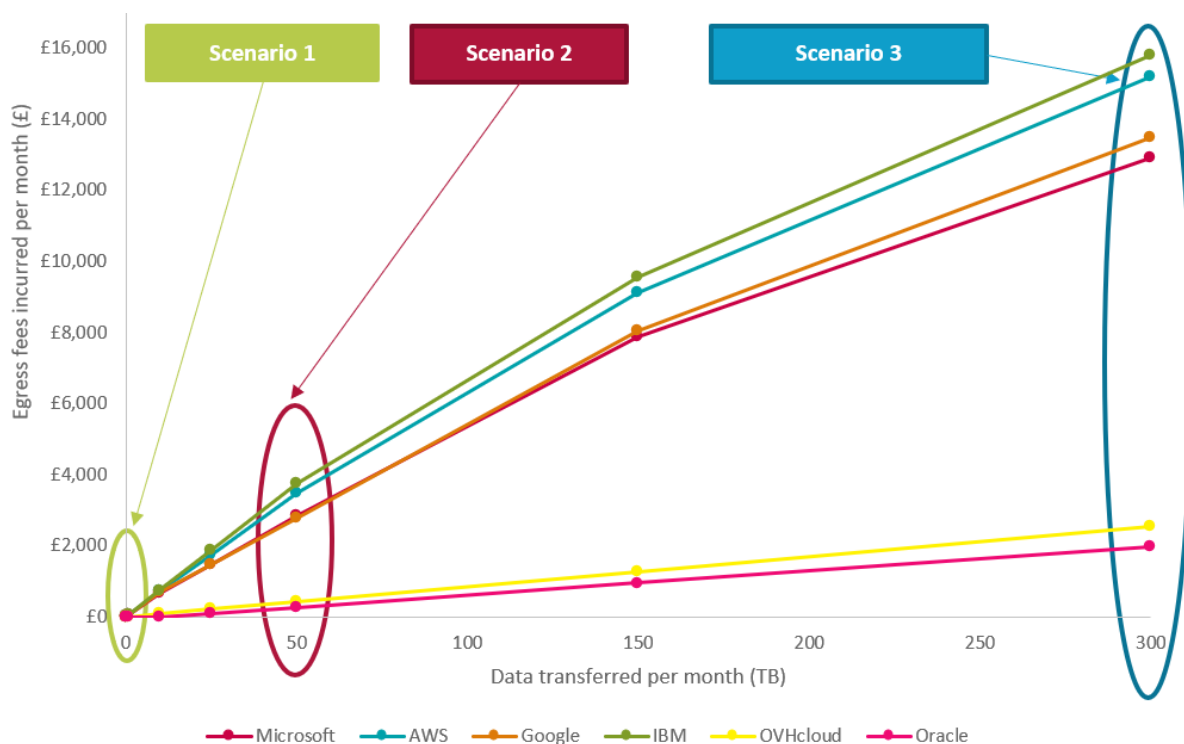
A company collects event data from mobile devices. It receives and processes around 1.5TB per day, meaning it receives around 50TB per month. It has its servers and storage on one cloud but wants to use the analytics service of another cloud. To do so from AWS, Google or Microsoft would incur egress fees of around £2,700-£3,500 per month.

Scenario 3: High data transfer – moving data from storage for video analytics

A company is offering a video analytics service which is very CPU/GPU intensive. It handles around 10TB of video per day, around 300TB per month. If it was to ingest and store its data on either AWS, Google or Microsoft's cloud, and transfer it to the compute services of another cloud, it would incur egress fees of £12,500-£15,500 per month.

³⁸⁵ This level of data transfer would be captured within Oracle's free tier.

Figure 5.14: Egress fees per month for stylised example multi-cloud use cases



Source: Ofcom.

Egress fees can lead to a large increase in the effective price of a service

- 5.98 Egress fees may significantly increase the cost of taking a service from a different cloud provider. As an illustrative example, consider a customer who stores its data in AWS, and is considering whether to undertake its data analytics using AWS’s service (using AWS Redshift) or in Google (using BigQuery).
- 5.99 Taking the specific example of Scenario 2 from Box 5.13 above (i.e. analysing a dataset of 50TB), using BigQuery would incur a cost of £2,350 per month³⁸⁶ to analyse the data and around £3,500 per month in egress fees for moving 50TB of data from AWS to Google. In other words, the egress fees would more than double the effective price paid by the customer to analyse the data on Google’s cloud.³⁸⁷ Unless AWS’s Redshift product is more than twice as expensive as running the analytics on Google BigQuery, it would be materially cheaper for the customer to store and analyse the data on AWS’s platform. Similarly, if the customer was storing 50TB of data in Microsoft, it would need to pay egress fees of around £2,800 per month to move this data to Google. In both of these

³⁸⁶ The pay-as-you-go price of Google BigQuery is no more than £0.04 per GB to store data within BigQuery, and a further £0.005 per GB for on-demand analysis of that data. Which would imply that at 50TB worth of data, the use of BigQuery would incur active physical storage costs around £2,100 per month and the cost of analysing that data would be around £250 per month. This is likely to be an overestimate of the price of using BigQuery in this scenario, as it assumes all 50TB of data is stored in active physical storage from day 1 of the month, rather than transferred in smaller flows throughout the month.

³⁸⁷ In this example we abstract from any internal data transfer charges which a customer may face when moving its data within AWS cloud, but, in any case, these are lower than egress fees for transferring data out via the internet.

cases, it is unlikely that Google would be able to afford to provide discounts sufficient to offset these egress fees.

Egress fees can be a material proportion of actual spend for many customers

5.100 A customer's spend on cloud services will vary based on a number of factors. However, we understand that 70% of [redacted] pay less than £1,000 per month.³⁸⁸ For such customers, the egress fees incurred in Scenario 2 outlined in Box 5.13 above (in the region of £2,700-£3,500 per month) - would be significantly more than the average monthly spend. Even for a customer spending £1 million per annum, the high data transfer scenario could lead to egress fees representing a 20% addition to total cloud spend – and we understand that most UK cloud customers spend a lot less than £1 million per annum on cloud services.

Egress fees can create unpredictability for customers

5.101 We also understand that egress fees can create unpredictability for customers. This is because it can be difficult for customers to predict the nature of their cloud usage over time. This additional uncertainty would be lower if all services were purchased on the same cloud (because the customer would not be liable to egress fees).³⁸⁹ This can act as a further barrier to purchasing services from multiple clouds by making it difficult to compare prices between different providers. We explore how this, and other price transparency concerns affect competition in more detail in paragraphs 5.167-5.177.

Egress fees can also impact switching

5.102 As set above, egress fees create a one-off switching cost for customers when looking to move their storage data from one provider to another. In Box 5.15 below, we use the scenarios from Box 5.13 above to illustrate the potential egress fees a customer would incur when switching its data from one cloud provider's storage service to another. This will vary based on the data the customer has stored in the cloud, which will be a function of the data intensity of the customer use case, and the length of time the customer has been ingesting data to the cloud. We have also heard that egress fees when switching could be particularly large in relation to the transfer of data warehouses.³⁹⁰

Box 5.15: Egress fees and switching

We extend the scenarios from Box 5.13 above to illustrate the potential egress fees a customer would incur when switching its data from one cloud provider's storage service to another.

Scenario 1: Low data storage -> report generation

³⁸⁸ [redacted]. Based on figures from [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

³⁸⁹ The customer may still face internal data transfer charges if using a single cloud for all its needs – but these are lower than egress fees.

³⁹⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

The company generating reports is storing 300GB of reports per month. After two years, it will have generated around 7TB of data. The company seeks to transfer its data to another cloud provider. This would incur egress fees of around £500.

Scenario 2: Medium data storage -> moving data from storage to database

The company analysing metric data is storing around 50TB per month. After two years, it will have generated around 1PB of data. Switching this data to another cloud provider would cost £35,000-£45,000.³⁹¹

Scenario 3: High data storage -> moving data from storage for video analytics

The video analytics company handles 300TB per month. Over three years, this will have created around 10PB of data. Switching this this data to another cloud provider would cost £330,000 - £420,000.

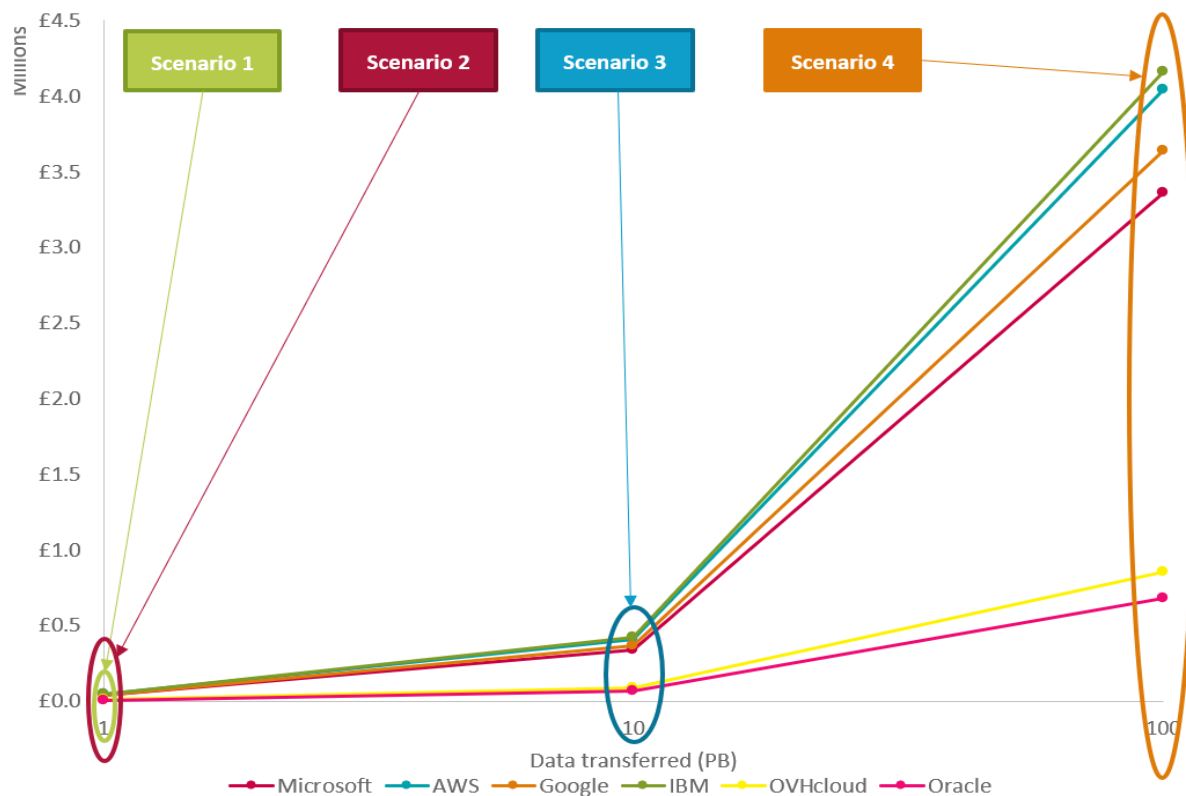
Scenario 4: Very high data storage

We understand that some customers have much higher volumes of data stored in the cloud. For example, we understand that [redacted] [20-30] customers, including [redacted] all have in excess of 100PB stored in AWS S3.³⁹² Switching this amount of data to an alternative provider would cost millions of pounds in egress fees.

³⁹¹ We understand that [redacted] has around 700TB in cloud storage – and so it would face egress fees of this order of magnitude were it to switch its data to a different cloud storage provider. [redacted] response dated [redacted] to the s.174 notice dated [redacted], page [redacted].

³⁹² [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

Figure 5.16: Egress fees for stylised switching scenarios



Source: Ofcom.

- 5.103 While these potential egress fees are larger in magnitude than those associated with transferring data as part of a multi-cloud architecture, they will only be occurred on a one-off basis when that customer seeks to switch its services. We have heard from some customers that, at present, the financial cost of egress fees were they to switch cloud providers could be at least £100,000 - although this may be a relatively small proportion of their current cloud spending ([redacted]%).³⁹³
- 5.104 When a customer is switching from one provider to another, it may be possible for a competing cloud provider to give the customer a discount to offset the one-off cost of paying egress. 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. This is more likely to be the case where the total volume of data being transferred is low. We also note that there may be additional costs involved in switching from one cloud provider to another, other than egress. This may subdue the impact of egress fees on switching at present.
- 5.105 We also understand that for very large data transfers, the customer may seek alternative ways of moving data (for example, using a private connection, or using a content delivery

³⁹³ See [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

network when broadcasting and delivering web content). This would lead to lower egress fees than transferring data via the public internet, but is still likely to be expensive.

- 5.106 We expect data volumes per customer to continue to grow over time, as existing users migrate more workloads to the cloud.³⁹⁴ The growth of rich media content and unstructured data³⁹⁵ is likely to increase the amount of data being stored, processed and analysed in the cloud.³⁹⁶ The egress fees associated with switching all of a customer's data from one storage service to another could therefore rise over time, and so increase the impact of egress fees on customer decisions on using a multi-cloud architecture or to switch.

Free tiers allow some data to be moved across clouds without charge, but are unlikely to be sufficient for all customers

- 5.107 Both AWS and Microsoft include free tiers, allowing 100GB per month of data to be transferred out of their cloud without charge. AWS provided information to show that its free tier (100GB per month) means that a large number of customers [X]% do not have to pay any fees for transferring data out of AWS, because these customers do not transfer more than the free tier allowance.³⁹⁷ In contrast, we note that over each month of 2021, 50-65% of Microsoft's active customers with egress usage at some point in 2021 did not pay any egress fees during the month because they did not meet the free monthly tier allowance, and that over the full calendar year 2021, only [X] [10-20]% of Microsoft Azure active cloud customers with egress usage did not pay any egress fees, because they did not in any month transfer more than the free monthly tier allowance – with [X] [60-70]% of active customers incurring egress fees at least once during the year.³⁹⁸
- 5.108 Observing that customers do not use more than 100GB per month does not necessarily imply that this is sufficient for customer needs. As explained above, we have seen evidence of customers adjusting their architecture to avoid paying egress fees. It is possible that customers may similarly adjust their service choice to use less than 100GB of data egress per month to avoid paying egress fees. In any case, a 100GB free tier is unlikely to be sufficient for all customers. In particular, free tiers are less likely to be sufficient for customers seeking to switch their data to another provider, where higher volumes are more likely to be involved.

³⁹⁴ The market research found that 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. See Context Consulting research report, slide 43; Section 3, paragraph 3.34 and Figure 3.4.

³⁹⁵ For example, a UHD video file is generally 4x the size of HD files.

³⁹⁶ For example, see [26 Cloud Computing Statistics, Facts & Trends for 2023 \(cloudwards.net\)](#) [accessed 14 March 2023].

³⁹⁷ AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, paragraph 11.7; and AWS response dated 17 February 2023 to our follow-up email dated 27 January 2023 concerning the s.174 notice dated 24 October 2022, paragraphs 12.3 -12.6, tables 8, 9 and 10.

³⁹⁸ Customers with annual customer revenue of greater than \$100. Microsoft response dated 18 February 2023 to follow-up email of 27 January 2023 concerning s.174 request of 21 October 2022, Q3.

Egress fees appear only loosely related to cost

- 5.109 We have also considered the relationship between egress fees being charged by providers and the costs of providing egress. We have not undertaken a detailed cost modelling assessment, but we have identified a number of indicators that egress fees are being set above the incremental cost of providing data transfer – namely:
- The difference in the prices charged by the hyperscalers and other providers such as OVHcloud and Oracle;
 - The prices charged by the hyperscalers are much higher than the likely transit charges they face, which we consider to be the key cost unique to data transfers out; and
 - Other aspects of the structure of egress fees charged by the hyperscalers, including use of free tiers.
- 5.110 Many smaller cloud providers offer much lower egress fees than the hyperscalers. As shown in Figure 5.10 above, AWS, Google and Microsoft charge fees which are 5-10 times higher than some smaller cloud providers, such as OVHcloud and Oracle. These smaller cloud providers could face higher costs of providing egress due to their smaller size, because they may not benefit from the same discounts as larger providers when negotiating with transit suppliers. We also understand that a group of smaller cloud providers has agreed to waive egress fees when transferring data between each other.³⁹⁹ Both Microsoft and Google offer discounted (although not zero) fees as part of this arrangement.⁴⁰⁰ This is consistent with the hyperscalers charging fees which are above their incremental costs.
- 5.111 Many cloud providers argue that the key cost component of external data transfer is transit.⁴⁰¹ This cost will depend on the transit and peering arrangements they have in place for exchanging traffic with other networks.⁴⁰² Transit services will incur charges. We understand that transit prices have fallen over time.⁴⁰³ However, the price of egress has fallen very little over the same period, implying that prices have diverged from cost over time.⁴⁰⁴

³⁹⁹ For more detail see [Bandwidth Alliance | Reduce Data Transfer Fees | Cloudflare](#).

⁴⁰⁰ AWS is not part of the Bandwidth Alliance.

⁴⁰¹ See OVHcloud response dated 17 November 2022 to the s.174 notice dated 27 October 2022, page 26; [redacted] response dated [redacted] to the s.174 notice dated [redacted], part B paragraph [redacted], Cloudflare, [AWS's Egregious Egress \(cloudflare.com\)](#), [accessed 14 February 2023], [redacted] response dated [redacted] to the s.174 notice dated [redacted], part B, question [redacted], page [redacted].

⁴⁰² These can include transit via the public internet, Private Peering (a direct private connection between the operators) or Public Peering (routing traffic via an Internet exchange point). BT Wholesale website. [Wholesale Internet Connect](#) [accessed 14 February 2023].

⁴⁰³ Cloudflare finds that that the cost of bandwidth has fallen by 93% in the last decade. Cloudflare, [AWS's Egregious Egress \(cloudflare.com\)](#) [accessed 14 February 2023]. This is consistent with evidence from BT. See BT response to the CFI, page 9.

⁴⁰⁴ For example, AWS increased its free egress tier from 1GB per month to 100GB per month on 1 December 2021. However, AWS egress fees for volumes between 100GB and 50TB per month have remained unchanged since 1 December 2014, and egress fees for data transfers above 50TB have remained unchanged since 1 July 2011. 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 14 March 2023]. We note that [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], paragraph [redacted].

- 5.112 Further, we understand that many peering arrangements are settlement-free, and so will have almost no cost to the cloud provider once an agreement has been reached.⁴⁰⁵ [X] estimates that around [X]% of its global data transfers out to internet are sent over settlement-free peering, and Internet exchanges.⁴⁰⁶ We also understand that transit and paid peering arrangements are often billed at the 95th percentile, meaning that in a given month, the 36 hours with the highest traffic will not be billed for.⁴⁰⁷ This means that cloud providers will not be charged for the transit used for the largest data transfers. This implies that a significant proportion of external data transfers do not incur transit charges.
- 5.113 We therefore expect that the costs that the hyperscalers face in purchasing transit and peering services to be lower than the revenue they receive from egress fees. We understand that [X] paid £[X] for transit and peering arrangements ([X]) that physically interconnect with its infrastructure in the UK, while receiving [X] in revenue for external data transfer in the UK.⁴⁰⁸ [X] has explained that there are various other cost components that are also relevant to data transfers out of [X], as explained below.
- 5.114 Other providers (including the hyperscalers) argue that the costs of providing data transfer out include a broader set of costs relating to the investment in the construction, maintenance and improvement of their shared network infrastructure, including data centres.⁴⁰⁹
- 5.115 [X] provided us with estimates of its total revenue from global data transfer (£[X])⁴¹⁰ and its total cost of global data transfer of £[X] (around [X]% of its total cloud costs) – including both internal data transfers and external data transfers, for 2021.⁴¹¹ [X].
- 5.116 [X] explained that there are various cost components that it considered relevant to data transfers out. However, some of these costs do not appear to be relevant to the transfer of data, and some of these costs do not appear to be uniquely attributable to the transfer of data. [X]% (£[X]) relate to categories which it states are not directly associated with data transfer and are allocated to data transfer on revenue proportion ([X]).⁴¹² [X]. Further, many of these costs appear to us to be common to a number of different services offered by the cloud providers (for example, including the storage of data), and the extent to which they are uniquely attributable to data transfer is unclear. For example, [X]% of this cost

⁴⁰⁵ See for example, NSRC, [Transit and Peering](#) [accessed 14 February 2023] and Cloudflare response to the CFI, page 1.

⁴⁰⁶ [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X], paragraph [X].

⁴⁰⁷ See [WIK Report \(sipotra.it\)](#), section 2.3, p.38 [accessed 6 March 2023], [What is 95th Percentile \(Burstable\) Billing? | Blog | deeserve](#) and [The Relative Cost of Bandwidth Around the World \(cloudflare.com\)](#) [accessed 21 March 2023].

⁴⁰⁸ [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X], paragraph [X]; and [X] response dated [X] to the s.174 notice dated [X], question [X]. [X] also notes that data transfers out to the internet are moved from the UK to another region before exiting its network – and states that the transit and peering [X] costs associated with this routing should be included. Even this would only increase the total transit and peering [X] costs to around £[X]m – less than [X]% of the revenue received for these transfers.

⁴⁰⁹ [X] response dated [X] to the s.174 notice dated [X], page [X]; [X] response dated [X] to the s.174 notice dated [X], part B, paragraphs [X]; [X] response dated [X] to the s.174 notice dated [X], part B, question [X].

⁴¹⁰ See [X] response dated [X] to the s.174 notice dated [X], [X].

⁴¹¹ See [X] response dated [X] to the s.174 notice dated [X], part B, [X].

⁴¹² See [X] response dated [X] to the s.174 notice dated [X], part B, [X]; and [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X], paragraph [X].

(£[X]) relates to data centre costs.⁴¹³ The remaining [X]% (£[X]) relates to network expenses. This includes some cost categories which are directly and solely related to data transfers out ([X]), some categories of cost which relate to all internal as well as external data transfers ([X]) and some cost categories which have been scaled on the basis of bandwidth demanded (network equipment depreciation and maintenance).⁴¹⁴

- 5.117 Therefore, [X], we consider that many of the cost categories cited above are not uniquely attributable to data transfer, and the actual incremental cost of providing external data transfer is likely to be well below this cited total cost [X].⁴¹⁵ We plan to do further analysis on this issue ahead of our final report, and will continue to gather evidence to inform this.
- 5.118 Other elements of the structure of AWS, Google and Microsoft's data transfer prices imply that providers have an element of flexibility over how they charge for egress relative to cost. AWS and Microsoft include a free egress tier – allowing customers to egress a certain amount of data per month for free. Providers do not charge for ingress.⁴¹⁶ Indeed, some providers responded that charging for egress and not ingress was based on the industry standard rather than purely cost allocation.⁴¹⁷
- 5.119 The difference in prices for ingress and egress is likely to reflect the competition incentives faced by cloud providers. Historically, competition between providers has led to ingress fees being set at zero. 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.⁴¹⁸
- 5.120 However, competition between providers has not led to a significant reduction in the hyperscalers' egress fees and is unlikely to do so in the future. As shown in Figure 5.10 above, the three hyperscalers have broadly similar prices for egress. Very large customers may be able to negotiate fixed fees for transferring data out of the cloud.⁴¹⁹ However, in

⁴¹³ See [X] response dated [X] to the s.174 notice dated [X], part B, [X]; and [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X], paragraph [X]. We note that [X] suggests that data centre network costs are only relevant to internal data transfers, the costs of which are recovered through storage and compute charges. See [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X], question [X], page [X].

⁴¹⁴ See [X] response dated [X] to the s.174 notice dated [X], part B, [X]; and [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X], paragraph [X].

⁴¹⁵ [X] also argued that data transfers out create additional costs to cover the risks associated with not controlling the end-to-end infrastructure. However, we do not consider that these would materially increase the cost of providing data transfers out. [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X], paragraph [X].

⁴¹⁶ A cloud provider stated that the costs for a cloud provider for ingress and egress are fairly symmetric. [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X], question [X]. Another cloud provider said that the cost of using bandwidth is invoiced to the party that initiates the transfer (not the receiving party), which explains that ingress is almost always free. See OVHcloud response dated 17 November 2022 to s.174 notice of 27 October 2022, Q9, p.26.

⁴¹⁷ See [X] response dated [X] to the s.174 notice dated [X], question [X]; and [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X], question [X]. Also see [X] response dated [X] to our follow-up email dated [X] concerning the s.174 notice dated [X].

⁴¹⁸ ACM, May 2022, [Market Study Cloud Services](#), page 58.

⁴¹⁹ For example, [X] says it has negotiated fixed fees for data egress. [X] response dated [X] to the s.174 notice dated [X], question [X].

general, providers have weak incentives to attract customers to place data in their cloud by making it easy to move it out again, and so in the process making it more attractive for users to purchase services from other providers.

Provisional conclusion

- 5.121 We therefore provisionally conclude that there are strong indications that egress fees act as a barrier to customers from adopting multi-cloud architectures for workloads requiring data exchange. 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 relative to a range of customer spending levels that are representative of a large fraction of cloud infrastructure customers. This has the potential to impact the choices of a material fraction of customers, with many identifying egress fees as an important barrier to multi-cloud. We also observe that a significant fraction of respondents in the Context Consulting market research use data management and analytics services, and these services represent a material fraction of hyperscaler revenues.
- 5.122 We provisionally conclude that egress fees may impact customers incentives to switch between cloud providers – although we find that this impact is likely to be lower than the impact of egress fees on the ability to adopt multi-cloud architectures. Nonetheless, the costs can be material for those customers hosting large volumes of data in the cloud (e.g. in data warehouses), and nearly a quarter of respondents to the market research highlighted egress fees as a barrier to switching.
- 5.123 We also provisionally conclude that the egress fees charged by the hyperscalers are likely to be only loosely related to the cost of providing egress, and that the incremental cost of providing egress may be well below the prices currently being charged by the hyperscalers.
- 5.124 We plan to do further analysis on this issue ahead of our final report and will continue to gather evidence to inform this.

Committed spend discounts

- 5.125 Discounts are generally a positive feature of markets, leading to lower prices and indicating competition between providers on price. In Section 4 we note how cloud providers, in particular the hyperscalers, are using discounts to compete with each other to bring customers into their ecosystems. This subsection focuses on assessing the competitive impact of discounts on cloud infrastructure competition, and how they may contribute to customers' incentives to concentrate all or most of their workloads with one cloud provider.
- 5.126 Our assessment of discounts is focused on committed spend discounts, one of the types of discounts offered by cloud providers including the hyperscalers. Committed spend discounts involve a customer agreeing to spend a set amount on a cloud provider's products over a set period of time. In turn, the customer typically receives a percentage discount. Such discounts are commonly used by the hyperscalers, especially in negotiations

with the largest cloud customers – a large proportion of whom have significant spending commitments to the hyperscalers.

Use of committed spend discounts by cloud providers

- 5.127 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.⁴²⁰ All three hyperscalers also offer broad committed spend discounts, where spending across all (or almost all) of their cloud products will count towards drawing down the commitment, but only as part of privately negotiated agreements with large customers.⁴²¹ Our assessment focuses on these broad, privately negotiated committed spend discounts offered by the hyperscalers.
- 5.128 Committed spend discounts are offered alongside other types of discounts, such as for pre-paying for cloud services or credits for trialling new products.⁴²² The discount obtained through committed spend may apply to only the spending which falls within the commitment or to all of a customer's spending, including the portion above the commitment threshold.⁴²³ Customers making larger commitments will receive larger discounts in percentage terms.⁴²⁴
- 5.129 Generally, such commitments function as minimum spends, where a customer which does not reach their commitment pays the difference between their actual and committed spend.⁴²⁵ 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 commitment period.⁴²⁶
- 5.130 Spend commitments and associated discounts are often negotiated alongside other contractual provisions such as credits and services including technical support. Customers may also renegotiate agreements involving spend commitments which are in progress, particularly when they intend to increase their spending with a cloud provider beyond their original forecast and wish to negotiate a higher discount along with this.⁴²⁷
- 5.131 The broad committed spend discounts offered by the hyperscalers generally apply to customer spending on ISVs' products purchased through the hyperscalers' marketplaces,

⁴²⁰ See, for example, Google's documentation of its resource-based committed use discounts. Google website. [Resource-based committed use discounts](#) [accessed 7 February 2023].

⁴²¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]. [redacted] response to the s.174 notice dated [redacted], part A question [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

⁴²² [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

⁴²³ Hyperscalers indicated that the exact mechanisms of how discounts are applied to customers' spending varies from agreement to agreement as they are privately negotiated. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted].

⁴²⁴ [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted].

⁴²⁵ One hyperscaler told us that contractual damages may apply in some cases. [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]. [redacted] response to the s.174 notice dated [redacted], part A question [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

⁴²⁶ [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

⁴²⁷ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

as well as to spending on the hyperscalers' own cloud products. 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.⁴²⁸

- 5.132 As the broad committed spend discounts of the hyperscalers are individually and privately negotiated, there is little information publicly available on their size or how they scale with increased commitment. The hyperscalers emphasised that customer prices and discounts depend on a range of factors and are individual to each customer, though they said prices have trended down over time.⁴²⁹ The hyperscalers who were able to give us information on the sizes of their discounts told us that the bulk of their discounts are between [redacted]% relative to their pay-as-you-go prices, with only a very small percentage of discounts exceeding [redacted]%.⁴³⁰
- 5.133 Agreements between customers and the hyperscalers vary in duration, but those involving broad committed spend discounts are generally at least a year long. Agreements of longer than a year may specify individual thresholds within each year, with varying allowances to shift commitment draw down between years. We have seen examples of committed spend agreements with durations between one and seven years.⁴³¹ The hyperscalers told us that in 2022 [redacted]% of their total UK cloud revenues were accounted for by customers with spending commitments and agreements lasting at least one year, [redacted]% of revenues were accounted for by customers with spending commitments and agreements lasting at least 3 years, and [redacted]% 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 [redacted]% of revenues to [redacted]% of revenues.⁴³²
- 5.134 Committed spend discounts are one of the primary ways in which large customers are able to negotiate with the hyperscalers on price.⁴³³ The hyperscalers have told us that customers with privately negotiated committed spend discounts accounted for [redacted] of their UK cloud customers by count, but for [redacted]% of their UK cloud revenues. One

⁴²⁸ One of the hyperscalers which now uses a cap told us that they previously allowed marketplace spend to draw down commitments at half 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. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted].

⁴²⁹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

⁴³⁰ [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted].

⁴³¹ An example of a seven-year commitment to a hyperscaler was included in [redacted] response dated [redacted] to our customer questionnaire.

⁴³² One hyperscaler was only able to provide this information for 2021 rather than 2022. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted].

⁴³³ Financial Times, 1 January 2023. [Big Tech under pressure from cost-conscious cloud customers](#) [accessed 20 January 2023].

hyperscaler told us that [redacted] of their UK customers have such discounts in their agreements.⁴³⁴

How the hyperscalers' committed spend discounts can encourage purchasing most or all their needs from a single provider

- 5.135 These broad committed spend discounts may encourage customers to purchase most or all their needs within a cloud provider's ecosystem. This is a commercial factor that has the potential to encourage customers to place multiple workloads with the ecosystem of a single cloud provider (including both first-party products and third-party products purchased through providers' marketplaces) even where there are few technical barriers to using multiple cloud providers. This subsection examines the facts we have gathered on these discounts and assesses how they may affect the decisions facing customers.
- 5.136 We have limited evidence at this stage on exactly how these discounts are set, given that these are privately negotiated and concluded on a bilateral basis with customers. From the available evidence we do have, we provisionally think that the way in which the discounts are structured incentivises customers to purchase all or most of their needs from within a single hyperscaler's ecosystem.
- a) The hyperscalers told us that the percentage discount a customer receives across their entire spend increases as the amount they commit increases. This can provide an incentive for customers to place their spend with a single cloud provider if there is little else to distinguish the offerings of cloud providers.
 - b) Increases in the percentage discount received by customers as they increase their commitment amount happen in increments. This may create a distortion in a customer's choices when they are close to a threshold in a cloud provider's schedule of discounts and commitments. Under these circumstances, a small decrease in the customer's commitment to the cloud provider may cause them to fall beneath a threshold and decrease their discount across all of their remaining use. The customer may reject an alternative cloud provider's product at this margin, even if its offering is superior on its own merits, to avoid this fall in their discount. The more aggressive the discount structure – i.e. the larger the increase in the discount as customers pass spend thresholds – the more likely it is that the pricing structure could encourage customers to single-source. We currently have limited evidence on the size of the increments in percentage discounts used by cloud providers, and more generally on the relationship between commitment thresholds and the percentage discounts offered by the hyperscalers.
 - c) Private negotiation between the hyperscalers and customers to establish the discounts and thresholds may increase the risk that committed spend discounts contribute to a customer purchasing all or most of their needs from a single ecosystem. By tailoring the schedule of discounts to customers' product needs and total spending, the

⁴³⁴ [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted].

hyperscalers may have an improved ability to structure discounts to encourage single sourcing. A customer told us that the secrecy of the discounts used by the hyperscalers prevents them from determining how the discounts they negotiate compare to other customers.⁴³⁵

- d) Long durations of agreements may compound the effects of committed spend discounts by tying customers' spending into that provider until the agreement expires. Where many large customers have agreements with long durations, the amount of spending at any given time that rival cloud providers are able to compete for and win will be restricted. A large amount of customer spending is committed for long periods of time, with [redacted]% of the hyperscalers' total cloud revenue accounted for by customers which have agreed to contracts including spending commitments of 3 years or more.

- 5.137 Overall, the pricing structure created by the hyperscalers' use of committed spend discounts may create an incentive for customers to commit all or most of their spend with a single primary cloud provider. That said, a customer may still choose to use a cloud provider outside of their primary cloud provider's ecosystem, for example because they offer a high-quality product which would incorporate well into their cloud solution. Competing for parts of a customer's demand may be more challenging for the secondary cloud provider if the customer would receive a lower discount across their entire commitment with their primary cloud provider.

Evidence on how discounts impact customer choices

- 5.138 We have seen some evidence that the availability and use of committed spend discounts is affecting customers' decisions when purchasing cloud services. This evidence is mixed on whether the discounts may be encouraging customers to focus their spending within a single cloud provider's ecosystem, indicating that any impact is likely to depend on the circumstances of individual customers.
- 5.139 Some customers told us that they viewed the incentives offered by cloud providers as encouraging them to purchase multiple cloud services from the same cloud provider.⁴³⁶ Several customers said they make commercial decisions, for example on whether to use cloud marketplaces to purchase products, on the basis of being able to draw down their commitments.⁴³⁷ This demonstrates that these customers' commitments are set at a sufficiently high level that reaching these commitments is a commercial consideration for them, and that this is inducing some customers to purchase incremental needs within the same cloud provider's ecosystem.
- 5.140 This said, another customer told us that making spend commitments with their provider does not limit their choices around which products they choose. They stated that they are able to choose commitments that match their forecasted spend with each cloud provider

⁴³⁵ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁴³⁶ For example, [redacted] response dated [redacted] to our customer questionnaire.

⁴³⁷ For example, [redacted] response dated [redacted] to our customer questionnaire, question [redacted].

based on which products are most appropriate, meaning they are not left in the position of choosing products on the basis of being able to draw down spend commitments.⁴³⁸

- 5.141 Evidence gathered from a hyperscaler indicates that it is aware of the potential of committed spend discounts to induce customer retention and this is part of its strategy in using them. Internally it discusses the productivity of its enterprise cloud sales teams, for which obtaining customer commitments is a key metric. It notes that customers entering into commitments can be a springboard for increased adoption and can drive downstream increased sales productivity.⁴³⁹ The hyperscaler's strategy in negotiating commitments with larger customers is described as being to win workloads that disproportionately improve customer retention.⁴⁴⁰
- 5.142 These considerations may indicate that customers' commitments to the hyperscalers are sometimes being set at levels which could provide customers with an incentive to concentrate their spending on a single primary cloud provider. That said, our evidence on customer purchasing decisions shows that many customers use multiple cloud providers. This evidence is presented and discussed in Annex 7. This indicates that secondary cloud providers are able to overcome any barriers which committed spend on its own may create to multi-cloud, at least in relation to some customers and workloads.

Competition for customers' commitments

- 5.143 Cloud providers competing for customers' business will set commitment discounts alongside pay-as-you-go prices and other types of discounts. In some respects, committed spend discounts are a part of normal price competition between cloud providers. "Competition to win purchase commitments made by customers" is how one hyperscaler characterised competition in the cloud market as functioning in the presence of committed spend discounts.⁴⁴¹
- 5.144 Prevalent discounting may indicate strong competition between cloud providers – that all three hyperscalers offer committed spend discounts suggests that competition for large customers is happening on the basis of commitments. Two of the hyperscalers explained that customers agreeing to spending commitments allows them to forecast future capacity requirements more accurately, and so to invest in cloud infrastructure more confidently. One hyperscaler further said that the benefits from these improved forecasts are passed on to customers through the discounts they receive in exchange for commitment.⁴⁴²

⁴³⁸ Ofcom/ [redacted] meeting, [redacted].

⁴³⁹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

⁴³⁹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted].

⁴⁴⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted].

⁴⁴¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]. [redacted] response to the s.174 notice dated [redacted], part A question [redacted].

⁴⁴¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

⁴⁴¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

- 5.145 The hyperscalers, with their wide ranges of products across the cloud stack, are likely to be able to compete on a level footing for a customer's spend commitments when they first migrate to the cloud, or where customers can easily switch their existing cloud products between cloud providers when renegotiating contracts. All three hyperscalers use committed spend discounts and appear able to counter the use of a discount by another hyperscaler with their own offerings so long as their product range is suitable for the customer's needs.
- 5.146 Cloud providers with narrower ranges of products may find themselves in a different position. These providers may be less able to offer large customers a solution for all or most of their product needs. Such cloud providers may be reliant on competing for incremental use cases of customers who have a large amount of their spending with the hyperscalers. They may therefore be placed at a disadvantage by any incentives given to customers to focus most or all their spending with a single cloud provider, raising barriers to entry and expansion.
- 5.147 Some smaller cloud providers have told us that the hyperscalers' use of committed spend discounts is harming competition in cloud infrastructure services by increasing customer lock-in. A cloud provider submitted that various attributes of the committed spend discounts used by the hyperscalers induce customers to move as many workloads to that provider as possible, leading to lock-in. The attributes they identify are the length of agreements, the broad ranges of products they apply to (rather than being related to specific use cases), and the high monetary thresholds of the commitments.⁴⁴³ Other smaller cloud providers listed spending commitments as a contributor to customer lock-in, and a factor which increases switching costs.⁴⁴⁴
- 5.148 We recognise that we have limited information currently on the decision making of customers who are being offered committed spend discounts by their providers, and of the pricing structures created by the hyperscalers' use of the discounts. We will continue to gather evidence and views on how committed spend discount structures are set and the extent to which this encourages sourcing from a single cloud provider.

Impact on ISVs is mixed because they are included in spending commitments to some extent

- 5.149 Most ISVs focus on providing a small set of products, which can be a small component of large customers' overall cloud solution. As such, they may be particularly vulnerable to customers being encouraged to use a single cloud provider. However, the policies of the hyperscalers which allow spending on ISVs' product through their marketplaces to draw down commitments, sometimes up to a cap, mean that the impact of committed spend discounts on ISVs may be lesser than for rival cloud providers.
- 5.150 ISVs whose products are available through the hyperscalers' marketplaces will not face the problem of competing with the hyperscalers to win customers' commitments, and they are able to win the business of customers who have already agreed to commitments with the

⁴⁴³ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁴⁴⁴ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

hyperscalers. This may be expected to result in a larger proportion of ISV sales being transacted through the hyperscalers' marketplaces.⁴⁴⁵ A survey of cloud customers and ISVs found that drawing down committed spend was the most common reason for customers to use the hyperscalers' marketplaces, while accessing customers' spend commitments was among the most important reasons for ISVs listing their products on marketplaces.⁴⁴⁶

- 5.151 However, some ISVs may be negatively affected by the hyperscalers' use of committed spend discounts in the same way that smaller cloud providers are. One SaaS ISV told us that committed spend discounts are part of a group of practices which the hyperscalers use to encourage customers to purchase a "packaged business application" of cloud products across the stack from a hyperscaler's ecosystem. Consequently, it said ISVs need to offer a "best of breed" alternative in order to give customers a compelling alternative to the offerings of the hyperscalers.⁴⁴⁷ Another ISV told us that spending on its products through marketplaces draws down customers' commitments at a slower rate than spending on competing first-party products offered by the hyperscalers, and that this negatively impacted its ability to compete with the hyperscalers.⁴⁴⁸
- 5.152 Overall, the effects of committed spend discounts on ISVs appear mixed. The policies of the hyperscalers' marketplaces and the degree to which they are managed in a way which is fair to ISVs and their products is an important factor in this, especially in the longer term.

Implications for new workloads for renegotiating customers who have limited ability to switch

- 5.153 A customer who is renegotiating a committed spend discount with a provider whose products they have already adopted may be in a different position to when they were first negotiating with the provider. A customer's cost of switching to a different cloud provider is likely to increase over time as they use and invest in a cloud provider's products. Even where strong alternative products are available for the products the customer is currently using, the customer may not be able to switch easily.
- 5.154 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 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 threat of losing this discount can create a strong incentive to purchase incremental requirements from the ecosystem of their existing provider.
- 5.155 We have heard from several customers that in the course of renegotiating their agreements, the hyperscaler they were using required them to increase the amount of

⁴⁴⁵ Some customers told us this is a primary reason for their use of marketplaces. See, for example, [redacted] response dated [redacted] to our customer questionnaire.

⁴⁴⁶ Tackle website. [State of Cloud Marketplaces 2022](#) [accessed 2 February 2023].

⁴⁴⁷ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁴⁴⁸ Ofcom/ [redacted] meeting, [redacted].

spend they had committed in order to continue receiving any discount from committed spend at all. These concerns were from three customers of one hyperscaler, and one customer who did not specify which hyperscaler their concern applied to. On the renegotiation of committed spend discounts, the hyperscaler which was mentioned by three customers ([X]) submitted that customers who maintain their spending levels generally will not face price rises following renegotiation.⁴⁴⁹

- 5.156 The first of these customers told us that each time they have renegotiated their agreement with [X], they have been required to increase their spend commitment in order to continue receiving a discount comparable to their previous agreement, regardless of whether they had a business need to increase cloud spend. They noted that this was particularly relevant when customers are using provider-specific products and so cannot move to an alternative easily.⁴⁵⁰
- 5.157 The second of these customers told us their experience of renegotiating with [X] has also been one of large increases in commitment amount. This customer noted that they have become locked in to [X] over time and that this has limited their ability to respond to price rises. The customer added that meeting the high commitments may require customers to migrate workloads from alternative providers to [X]. Further, the customer stated that they would like to have the choice to operate a multi-cloud strategy, for commercial and operational reasons, but are prevented from doing so by the commercial practices of [X].⁴⁵¹
- 5.158 The third of these customers told us that [X] had renegotiated with them based on their historical and projected spend levels, and that they were only able to choose the discount level that was aligned with this spend.⁴⁵²
- 5.159 The fourth customer told us that during the renegotiation of a contract with any hyperscaler, the negotiation power shifts markedly to the hyperscaler because switching is extremely difficult unless a customer uses exclusively basic or provider-agnostic services (especially IaaS) and/or have strict and complex utilisation restrictions. They told us that this means that a hyperscaler in a renegotiation is able to push for an increase in spend commitment compared to current spend in order for a user to maintain current benefits.⁴⁵³
- 5.160 Customers may be pushed by this type of conduct to reduce the extent to which they use multiple providers, even for new workloads which are relatively siloed and would be technically feasible to host on separate cloud infrastructure. As a consequence of customers facing high switching costs for some of their products, the hyperscalers may use this base of spending to induce customers to bring incremental spend into their commitments during renegotiation. This may make it harder for competitors to compete for incremental workloads which would otherwise be contestable.

⁴⁴⁹ [X] response dated [X] to the s.174 notice dated [X], part A question [X].

⁴⁵⁰ [X] response dated [X] to our customer questionnaire, questions [X].

⁴⁵¹ [X] response dated [X] to our customer questionnaire.

⁴⁵² [X] response dated [X] to the s.174 notice dated [X], question [X].

⁴⁵³ [X] response dated [X] to the s.174 notice dated [X], question [X].

Provisional conclusion

- 5.161 The use of committed spend discounts by all three hyperscalers reflects competition between them by allowing customers to use their commitments as a bargaining tool. Customers moving to the cloud can benefit from this type of competitive interaction.
- 5.162 The evidence we have seen at this stage indicates that the committed spend practices used by providers have the potential to encourage large customers to use a single hyperscaler for all or most of their cloud needs. This is because they can maximise the percentage discount they receive by concentrating their spend with a single provider.
- 5.163 Further, when a customer has new incremental workloads they wish to place on the cloud, they may prefer to add this spending to any existing commitment they have with their primary provider, as they will benefit from the discount on their new spending and will increase their discount across their existing spend. The threat of losing a discount on products which they cannot easily switch away from may be used by a hyperscaler to induce customers to greater amounts of spending within their ecosystem. This may dampen competition for customers' new workloads if this makes it difficult for rival cloud providers to compensate for the loss of a lower discount across their entire commitment by pricing low on the incremental spend alone.
- 5.164 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. 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 using only the narrower set of products they provide. They are more likely to rely on being the secondary provider of a customer and the use of committed spend discounts by hyperscalers may also limit their ability to compete for the incremental workloads of customers that may have otherwise been contestable.
- 5.165 Notionally, the hyperscalers may be able to contest for the incremental workloads by competing for the full commitment of a customer (for existing and new workloads). In principle, this could encourage competition between the hyperscalers by strengthening the bargaining position of customers. However, where the ability of customers to choose an alternative cloud provider for all their needs is limited, including due to high switching costs or a need for a specific provider's products, it may be more difficult for even larger rivals to challenge for such incremental workloads.
- 5.166 Due to the private and individualised nature of these discounts and of customers' cloud needs there remains a lot about the magnitude of their impact which is unclear. We would welcome more information on how spending commitments and discounts are set in negotiations, and how this affects the decisions of cloud customers.

Challenges predicting cloud spend

- 5.167 This subsection sets out our initial assessment of the impact of price transparency and complexity of usage and billing for cloud customers and the associated challenge some customers face when trying to predict future cloud spend.
- 5.168 In the quantitative part of the market research, value for money was rated as the second most important reason for picking a cloud provider. Microsoft's response to our CFI echoed this, saying that the cost of cloud solutions was a key consideration for customers. They also noted that it was important for customers to see how costs compared to other options, such as relying on their own hardware infrastructure.⁴⁵⁴
- 5.169 Findings from the 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.⁴⁵⁵ In addition, 52% of customers said they were concerned or very concerned about the lack of pricing transparency.⁴⁵⁶ Customers in the qualitative part of the market research told us that 'bills can be confusing'⁴⁵⁷ and they find that cloud providers are often 'trying to muddy the waters in terms of costing'.⁴⁵⁸
- 5.170 Other customers told us similar stories. BT'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.⁴⁵⁹ [redacted] 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.⁴⁶⁰
- 5.171 [redacted] 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.⁴⁶¹
- 5.172 We have also heard from some smaller cloud providers of the challenges they face when trying to give customers a competitive offer, given the lack of understanding of hyperscaler pricing. Hyve Managed Hosting told us that 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 the problem, it highlighted the complexity of the issue and noted a series of examples. Hyve Managed Hosting 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

⁴⁵⁴ Microsoft response to the CFI, page 6.

⁴⁵⁵ Context Consulting research report, slide 95.

⁴⁵⁶ Context Consulting research report, slide 131.

⁴⁵⁷ Context Consulting research report, slide 33.

⁴⁵⁸ Context Consulting research report, slide 58.

⁴⁵⁹ BT CFI response, page 25.

⁴⁶⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁴⁶¹ [redacted] response dated [redacted] to the s.174 dated [redacted], question [redacted].

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.⁴⁶² [X] 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.⁴⁶³

- 5.173 However, this is not an issue that every customer faces. We heard from various customers including [X], [X] and [X] that they find forecasting cloud infrastructure spending relatively issue-free,⁴⁶⁴ and they do not encounter ‘high volatility of prices or find pricing particularly opaque’.⁴⁶⁵
- 5.174 We have also seen some evidence of cloud providers awareness of the issues with transparency of billing. [X] documents showed that customers have told it that it is becoming increasingly complex to compare the true cost of [X] with competitors. It noted that this complexity comes from different discount plans and pricing structures compared to single discount structures elsewhere in the market.⁴⁶⁶
- 5.175 Customers in the qualitative part of the market research also told us how they had seen improvements in cost monitoring tools available from cloud providers.⁴⁶⁷ All three hyperscalers now have their own cost management tools available, AWS Cost Explorer, Azure Cost Management and GCP Billing, suggesting an intention to help resolve these issues. However, some customers explained that these tools are still often confusing, and there was a consensus among customers that vendors could do more to help manage their cloud spend.⁴⁶⁸

Provisional conclusion

- 5.176 Overall, the evidence indicates that in some circumstances customers may find it difficult to gain full visibility of usage prices and/or egress fees. 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 specific business practices such as the use of different pricing models across providers or a lack of clarity on the exact value of prices and fees.
- 5.177 Where customers find it difficult to predict and compare future spend when choosing a provider, we preliminary consider that this may prevent these customers 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

⁴⁶² Ofcom/CISPE/Hyve Managed Hosting meeting, 2 February 2023.

⁴⁶³ Ofcom/ [X] meeting, [X].

⁴⁶⁴ [X] response dated [X] to the s.174 notice dated [X], question [X]; [X] response dated [X] to the s.174 notice dated [X], question [X]; and [X] response dated [X] to the s.174 notice dated [X], question [X].

⁴⁶⁵ [X] response dated [X] to the s.174 notice dated [X], question [X].

⁴⁶⁶ [X] response dated [X] to the s.174 notice dated [X], [X].

⁴⁶⁷ Context Consulting research report, slide 80.

⁴⁶⁸ Context Consulting research report, slide 98.

compounds the barriers to entry and expansion for smaller cloud providers for whom price is a key lever to attract new customers.

Provisional conclusion on barriers to switching and multi-cloud

- 5.178 We provisionally conclude that the barriers we identify in this subsection can combine to substantially limit the ability of some customers to switch and to multi-cloud.
- Technical barriers increase the cost of any type of switching or multi-cloud. This is likely to pose the most material constraint when switching clouds requires porting a significant number of applications which are tightly integrated with proprietary services of the cloud provider. It is also likely to pose material limitations on integrated multi-cloud.
 - Egress fees are most likely to hinder the adoption of multi-cloud deployments which require substantial movement of data between workloads hosted on different clouds. They can also limit the likelihood of switching between cloud providers for those customers who would have to transfer large amounts of data.
 - Committed spend discounts encourage large customers to purchase all or most of their cloud needs from the ecosystem of their primary cloud provider, reducing their incentives to multi-cloud.
 - Challenges predicting cloud spend may prevent customers from exercising effective choice in selecting the cloud architecture that is most appropriate for their needs, which can reduce competition between cloud providers based on quality and price.
- 5.179 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 moved between cloud providers. However, our evidence suggests that a significant number of users are concerned about their lack of ability to switch and have had to settle for a sub-optimal approach to multi-cloud.
- 5.180 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.
- 5.181 Because of these barriers to switching and multi-cloud, 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 across the cloud layers 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 switching and multi-cloud.

5.182 As a result of these dynamics, 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, these barriers to switching and multi-cloud, some of which are exacerbated by specific business practices of the hyperscalers, have the potential to lessen effective competition in cloud infrastructure services:

- a) Cloud providers, particularly AWS and Microsoft, will face a weaker threat that some customers will switch or multi-cloud to make savings or purchase a rivals' best-in-breed solutions. It also reduces 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 providers' broad ecosystem, and this can impact a customers' existing and incremental workloads.
- b) Where smaller cloud providers are unable to compete for a customer's entire cloud needs, barriers to multi-cloud can further raise barriers to entry and expansion. 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.

Question 5.1: Do you agree with our analysis of potential barriers to switching and multi-cloud? As part of this:

- a) Please provide your views on to what extent and in what ways egress fees are a barrier to switching and multi-cloud. Please also provide your views on the extent to which egress fees currently charged relate to the incremental cost of providing egress.
- b) Please provide your views on whether specific business practices of cloud providers, particularly the hyperscalers, exacerbate technical barriers to switching and multi-cloud.
- c) Please provide your views on how commitment spend discounts are set and the impact these discounts have on the incentives of customers to multi-cloud.

Barriers to entry and expansion

5.183 In this subsection 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. We also explore the extent to which barriers to entry and expansion can be easier to overcome for cloud providers which are part of large tech conglomerates with major non-cloud businesses.

5.184 Competing as a vertically integrated cloud provider requires significant and on-going investment in several different areas:

- physical infrastructure, including data centres, servers and network equipment;
- a broad product portfolio (including solutions for specific industries); and
- customer relations.

- 5.185 The investments made in all three of these areas are largely fixed in nature, they create significant economies of scale, and 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.
- 5.186 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 the next subsection.

Investment in physical infrastructure

- 5.187 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.⁴⁶⁹
- 5.188 This subsection is structured as follows:
- a) Investment in data centres;
 - b) Economies of scale in data centres; and
 - c) Innovation in underlying hardware.

Investment in data centres

- 5.189 The hyperscalers have hundreds of data centres located around the world – for example, Microsoft Azure operates [X] [between 200-350], with plans to build more in the foreseeable future.^{470, 471} In comparison, smaller providers operate fewer data centres – for example, IBM operates [X] [between 50-100] and OVHcloud operates 33 data centres.⁴⁷²

Table 5.17: Data centre figures⁴⁷³

| | Global data centre count (owned and leased / co-located, 2021) | Average global data centre capacity (MW per data centre, 2021) |
|------------------------|--|--|
| AWS | [X] | [X] |
| Google | [X] | [X] |
| Microsoft Azure | [X] | [X] |

⁴⁶⁹ 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.

⁴⁷⁰ Microsoft response dated 9 December 2023 to the s.174 notice dated 21 October 2022, part B question 26a (Confidential Annex B26).

⁴⁷¹ Microsoft News, 2021. [Microsoft's virtual datacenter grounds 'the cloud' in reality](#) [accessed 28 March 2023].

⁴⁷² 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 24 March 2023].

⁴⁷³ We note that these providers may be using some of their data centre capacity for purposes other than providing cloud services, so these figures are only intended to be illustrative estimates.

| | Global data centre count (owned and leased / co-located, 2021) | Average global data centre capacity (MW per data centre, 2021) |
|----------|--|--|
| IBM | [REDACTED] | [REDACTED] |
| Oracle | [REDACTED] | [REDACTED] |
| OVHcloud | [REDACTED] | [REDACTED] |

Source: Ofcom analysis based on responses to s.174 notices.

- 5.190 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).⁴⁷⁴ A study by Emerson Network Power and Ponemon Institute estimated 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.⁴⁷⁵
- 5.191 A report by Dgtl Infra states that it could cost between \$7-12 million per megawatt of commissioned IT load to build a data centre.⁴⁷⁶ This means that build costs per data centre alone could be significant, in the tens to hundreds of millions. 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.
- 5.192 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 public cloud providers have entered the market in recent years by buying co-location space from providers of data centres, e.g. Civo⁴⁷⁷ and OVHcloud.⁴⁷⁸ [REDACTED].⁴⁷⁹ [REDACTED].⁴⁸⁰

⁴⁷⁴ Based on analysis of responses to our statutory information requests.

⁴⁷⁵ 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, IT assets, 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 16 March 2023].

⁴⁷⁶ 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 23 March 2023].

⁴⁷⁷ 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.

⁴⁷⁸ OVHcloud rented its first data centre in 2001, followed by a fully owned data centre two years later and today has 28 data centres in 8 countries around the globe. OVHcloud website. [The OVH Story](#) [accessed 17 February 2023].

⁴⁷⁹ [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], part B question [REDACTED]; [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], part B question [REDACTED]; [REDACTED] response dated [REDACTED] to our follow-up email dated [REDACTED] concerning the s.174 notice dated [REDACTED], part B question [REDACTED].

⁴⁸⁰ [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], part B question [REDACTED]; [REDACTED] response dated [REDACTED] to our follow-up email dated [REDACTED] concerning the s.174 notice dated [REDACTED], part B question [REDACTED]; [REDACTED] response dated [REDACTED] to our follow-up email dated [REDACTED] concerning the s.174 notice dated [REDACTED], part B question [REDACTED]; [REDACTED] response dated [REDACTED] to our s.174 notice dated [REDACTED]; [REDACTED] response dated [REDACTED] to our follow-up email dated [REDACTED] concerning the s.174 notice dated [REDACTED], part B question [REDACTED]; [REDACTED] response dated [REDACTED] to our follow-up email dated [REDACTED] concerning the s.174 notice dated [REDACTED], part B question [REDACTED]; and [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], part B question [REDACTED].

- 5.193 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 [redacted] 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.⁴⁸¹
- 5.194 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 – [redacted] 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, [redacted] suggests that limited access to real estate is an important barrier to entry and expansion.⁴⁸² Similarly, [redacted] also says that capacity can vary by location. For example, 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.⁴⁸³ [redacted] also highlights challenges regarding scarcity of available development land suitable for data centres – but notes that these capacity issues apply to building as much as leasing/co-location.⁴⁸⁴
 - b) Availability challenges in co-location data centre – 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. [redacted] says that if there is 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.⁴⁸⁵ Similarly, [redacted] notes 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.⁴⁸⁶

⁴⁸¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part B question [redacted].

⁴⁸² [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], part B question [redacted].

⁴⁸³ [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], part B question [redacted].

⁴⁸⁴ Ibid.

⁴⁸⁵ [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], part B question [redacted].

⁴⁸⁶ [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], part B question [redacted].

- 5.195 Our evidence suggests that 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) 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, [redacted] 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 the hyperscaler regions due to potential latency issues which would be critical for use-cases requiring a [redacted].⁴⁸⁷ [redacted].⁴⁸⁸
- 5.196 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 providers. We observe that small scale providers have much fewer data centres and are located in fewer geographic regions.⁴⁸⁹ Of the mid-scale 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 [redacted] [between 150-200] data centres globally).⁴⁹⁰ Oracle has also launched public cloud data centres in several geographic regions to rival AWS, Microsoft and Google.⁴⁹¹
- 5.197 Nevertheless, 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).⁴⁹² 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.

⁴⁸⁷ [redacted] response dated [redacted] to the s.174 dated [redacted], part A question [redacted].

⁴⁸⁸ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part B, question [redacted]; [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], part B question [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], part B question [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], part B question [redacted].

⁴⁸⁹ For example, OVHcloud has 33 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 23 March 2023]; Scaleway website. [Scaleway Datacenter homepage](#) [accessed 23 March 2023]; Civo website. [Frankfurt region now live](#) [accessed 23 March 2023].

⁴⁹⁰ Oracle response dated 9 December 2023 to the s.174 notice dated 31 October 2022, part B question 22a.

⁴⁹¹ Analysis of provider websites (accessed March 2022) suggests that AWS, Microsoft, Google and Oracle all offer public cloud services to customers from over 30 geographic regions located around the world. Also, based on 2021 data, IBM and Oracle both operate 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.

⁴⁹² For example, AWS offers Availability Zones (AZs) which are one or more discrete data centers 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 20 March 2023].

Economies of scale in data centres

- 5.198 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.
- 5.199 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 for all cost categories – but especially for energy and operating costs, which together account for 80% or more of annual data centre costs.⁴⁹³ Additionally, a Microsoft study states that investments in security and reliability also benefit from economies of scale, as these are largely fixed costs.⁴⁹⁴ There is also evidence that data centres with higher average rack density experience lower unit costs than data centres with lower rack density.⁴⁹⁵
- 5.200 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 5.17). For example, Microsoft’s data centres have an average capacity of [redacted] [10-20] MW per data centre globally compared to Oracle’s data centres which have an average capacity of [redacted] [less than 5] MW.⁴⁹⁶ When we compare infrastructure costs to revenue based on information provided by the hyperscalers, [redacted].⁴⁹⁷ [redacted], this trend is consistent with economies of scale, particularly in relation to data centre usage costs.
- 5.201 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.⁴⁹⁸ Whilst [redacted] told us that they expect the gap in server costs to be between 10-30% when compared to a hyperscaler, [redacted] 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.⁴⁹⁹ However, this may not be true during periods of shortages. [redacted] told us that in a context of

⁴⁹³ 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. Retrieved from [Wayback Machine](#) [accessed 16 March 2023]; results reported in Data Center Dynamics, 2016. [Research: Larger data centres make considerable savings on operating costs](#) [accessed 1 March 2023].

⁴⁹⁴ Microsoft, 2010. [The Economics of the Cloud](#) [accessed 26 January 2023].

⁴⁹⁵ Ibid.

⁴⁹⁶ 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.

⁴⁹⁷ [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], part B question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], part B question [redacted]; [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], part B question [redacted].

⁴⁹⁸ Microsoft, 2010. [The Economics of the Cloud](#) [accessed 26 January 2023].

⁴⁹⁹ Ofcom/[redacted] meeting, [redacted].

electronic components shortage, suppliers favour clients with larger purchasing volumes (i.e. hyperscalers) [§].⁵⁰⁰

- 5.202 There are also economies of scale associated with pooling of demand to reduce variance and improve the utilisation rate of servers and data centres. 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.⁵⁰¹ 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.⁵⁰² 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 explains 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.⁵⁰³
- 5.203 Our evidence on differences between the hyperscalers and smaller providers in terms of data centre/server utilisation rates is mixed. [§].^{504, 505}
- 5.204 Both Amazon 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.”⁵⁰⁶
 - 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.”⁵⁰⁷

Innovation in underlying hardware

- 5.205 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-

⁵⁰⁰ [§] response dated [§] to the s.174 notice dated [§], part A question [§].

⁵⁰¹ Microsoft, 2010. [The Economics of the Cloud](#), pp.5-6 [accessed 27 March 2023].

⁵⁰² 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](#), p.7 [accessed 16 March 2023].

⁵⁰³ IBM response dated 22 March 2023 to our proposed use of information dated 14 March 2023.

⁵⁰⁴ Provider responses to s.174 notices.

⁵⁰⁵ [§] response dated [§] to our follow-up email dated [§] concerning the s.174 notice dated [§], part B question [§].

⁵⁰⁶ Amazon, 2014. [Letter from CEO to shareholders](#), page 5 [accessed 1 March 2023].

⁵⁰⁷ Microsoft, 2022. [Microsoft Annual Report 2022](#) [accessed 27 March 2023].

based processors (Graviton3) provides up to 25% better compute performance and up to 60% less energy usage.⁵⁰⁸ Other examples, of hyperscalers customising hardware are provided in Box 5.18 below.

- 5.206 These efficiencies can be passed onto customers in terms of lower unit prices and increased performance.⁵⁰⁹ This in turn improves the price-performance ratio of the PaaS and SaaS services running on top of it. In other words, innovation in underlying hardware translate 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.⁵¹⁰ This is also supported by the responses we have received from large enterprises: [redacted], [redacted], [redacted], [redacted] and [redacted] all considered service quality to be an important factor when choosing cloud provider, with [redacted] and [redacted] ranking it as their most important consideration.⁵¹¹ [redacted] 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.”⁵¹²
- 5.207 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 (consequently benefiting from economies of scale and scope).⁵¹³ For example, Google’s TPU chips were initially used exclusively for its non-cloud services, e.g. Google Search.
- 5.208 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

⁵⁰⁸ 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 15 March 2023].

⁵⁰⁹ 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 15 March 2023].

⁵¹⁰ 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.

⁵¹¹ [redacted] response dated [redacted] to our customer questionnaire; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the our customer questionnaire; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁵¹² [redacted] response dated [redacted] to our customer questionnaire.

⁵¹³ 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 7 March 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”.⁵¹⁴
- b) Oracle has also been using Ampere CPUs and powerful graphics processors and GPUs from chipmaker Nvidia.⁵¹⁵ There are examples of Oracle using Ampere chips to improve its services and offer competitive prices.⁵¹⁶

5.209 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. 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 5.18: Examples of the hyperscalers customising hardware for data centres

AWS has customised hardware for its data centres, including i) Arm-based processors (Graviton CPUs),⁵¹⁷ ii) Nitro smart NICs/data processing units,⁵¹⁸ and iii) Inferentia⁵¹⁹ and Trainium⁵²⁰ ML accelerators. [redacted].⁵²¹ AWS launched its first Arm-based processor (Graviton CPU) in November 2018⁵²² and has since launched Graviton 2 and Graviton 3.⁵²³ 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).⁵²⁴

⁵¹⁴ Electronic Design, 2022. [Microsoft Taps Ampere’s Arm CPUs for New Cloud Service](#) [accessed 7 March 2023].

⁵¹⁵ “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 7 March 2023]. Furthermore, Oracle has made significant investments in Ampere. Protocol, 2022. [Oracle has pumped more than \\$400 million into chip startup Ampere](#) [accessed 7 March 2023].

⁵¹⁶ Oracle website. [Oracle Unlocks Power of Arm-based Processors at One Cent per Core Hour, Expanding Ecosystem, and Speeding App Development](#) [accessed 27 March 2023].

⁵¹⁷ AWS website. [AWS Graviton Processor](#) [accessed 6 March 2023].

⁵¹⁸ AWS website. [AWS Nitro System](#) [accessed 6 March 2023].

⁵¹⁹ AWS website. [AWS Inferentia](#) [accessed 6 March 2023].

⁵²⁰ AWS website. [AWS Trainium](#) [accessed 6 March 2023].

⁵²¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part B question [redacted].

⁵²² 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 7 March 2023]

⁵²³ AWS website. [AWS Graviton Processor](#) [accessed 7 March 2023].

⁵²⁴ 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 7 March 2023]; also see AnandTech, 2020. [Amazon makes Graviton2 AWS instances available](#) [accessed 7 March 2023].

It has been reported that Microsoft and Google will similarly customise their own Arm-based processors for their respective cloud services.⁵²⁵ [redacted].⁵²⁶ [redacted].⁵²⁷

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).⁵²⁸ 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 2015.⁵²⁹ 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.⁵³⁰ 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.⁵³¹

Benefits from sharing cloud infrastructure with non-cloud businesses

- 5.210 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.
- 5.211 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 thereby 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 2019.⁵³²
 - b) Microsoft uses Azure to power its Bing search engine, Xbox Live services and has migrated most of its Office365 services to Azure.⁵³³

⁵²⁵ Bloomberg, 2020. [Microsoft Designing Its Own Chips for Servers, Surface PCs](#) [accessed 7 March 2023]; and The Register, 2023. [Taking notes from AWS, Google prepares custom Arm server chips of its own](#) [accessed 7 March 2023].

⁵²⁶ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part B question [redacted].

⁵²⁷ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]

⁵²⁸ Google Cloud website. [Cloud TPU](#) [accessed 7 March 2023].

⁵²⁹ Google Cloud blog, 2017. [An in-depth look at Google's first Tensor Processing Unit \(TPU\)](#) [accessed 7 March 2023].

⁵³⁰ Ibid.

⁵³¹ Ibid.

⁵³² 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 14 March 2023].

⁵³³ ZDNET, 2021. [Microsoft moves closer to running all of its own services on Azure](#) [accessed 2 February 2022].

- c) Google announced plans to move parts of YouTube to its Google Cloud platform in 2021.⁵³⁴

5.212 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.

Provisional conclusion

- 5.213 We provisionally conclude that the need to invest in physical infrastructure is likely to create significant barriers to entry and expansion. 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 providers have entered the market and are increasing their physical infrastructure. However, they have significantly fewer data centres and in fewer regions.
- 5.214 We 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 providers having already made the capital investments and established expansive global networks of large data centres. Of the mid-scale 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.
- 5.215 Furthermore, the hyperscalers are likely to have an advantage due to the significant scale of their existing cloud customers (including an anchor tenant associated with their non-cloud business). Relative to other providers, they are likely to benefit from economies of scale in data centres given the size of their data centres and the ability to achieve higher utilisation rates in their data centres.⁵³⁵ The hyperscalers may also be better able to spread the fixed costs of their capital investments for further expansion.

⁵³⁴ DCD, 2021. [Google to migrate parts of YouTube to Google Cloud](#) [accessed 2 February 2022].

⁵³⁵ We do not have precise estimates of the global customer base for hyperscalers and other cloud providers, but in Annex 6 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

5.216 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.

The breadth and quality of providers' product ranges is an important parameter of competition

5.217 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.

5.218 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).⁵³⁶ 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.⁵³⁷ Furthermore, 84% of users have four use cases or more.⁵³⁸ In line with this, many of the customers responding to our customer questionnaire (e.g. [X]) noted they value range of services when picking a cloud provider.⁵³⁹ This is further reflected in customer purchase data, with the average hyperscaler customer purchasing multiple services.⁵⁴⁰

5.219 Product range is also important because a cloud provider with a broader product range will be able to attract a broader range of customers, including customers with specialist and/or narrow use cases, 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 ([X]), 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 service quality, reputation (including

⁵³⁶ Context Consulting research report, slide 69; and Context Consulting research data tables, Q25.

⁵³⁷ Context Consulting research report, slide 69; and Context Consulting research data tables, Q25.

⁵³⁸ Context Consulting research report, slide 35.

⁵³⁹ [X] response dated [X] to our customer questionnaire. Ofcom/[X] meeting, [X]. [X] response dated [X] to our customer questionnaire. [X] response dated [X] to our customer questionnaire. [X] response dated [X] to our customer questionnaire.

⁵⁴⁰ [X] response dated [X] to the s.174 notice dated [X], part B question [X]; [X] response dated [X] to the s.174 notice dated [X], part B question [X].

broad capabilities of AWS) and access to software only available on a specific cloud provider’s platform.⁵⁴¹

- 5.220 Customers also value their providers offering them access to the latest innovations – for example, suppliers of professional services and the hyperscalers have told us that ‘rate of innovation’ is an important capability customers consider when choosing cloud providers.⁵⁴² 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.⁵⁴³ It is therefore necessary for cloud providers to maintain a high rate of innovation to develop the range and quality of their products.
- 5.221 This explain why the hyperscalers have developed ecosystems of large product portfolios that span the full cloud stack, that allow customers to source products from one place and to easily combine them to build their IT solutions.
- 5.222 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”.⁵⁴⁴ Data we have gathered from providers (shown in Table 5.19 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 5.20 below indicates that AWS, Microsoft and Google are offering products across similar numbers of product categories.

Table 5.19: Number of cloud infrastructure products by provider – hyperscalers and mid-scale providers

| | Number of products |
|------------------|--------------------|
| AWS | 210+ |
| Microsoft | 200+ |
| Google | 190+ |
| Oracle | 110+ |
| IBM | 90+ |

Source: Ofcom analysis of IaaS and PaaS products listed on provider websites.

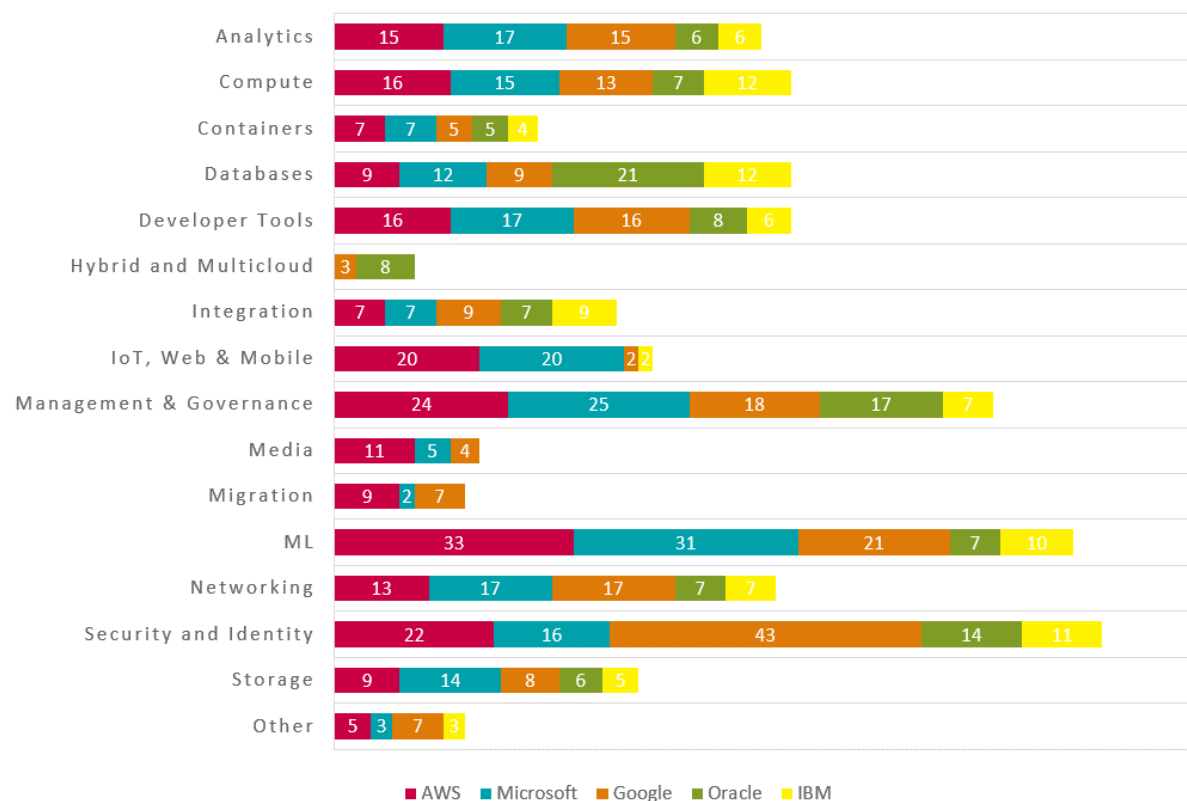
⁵⁴¹ [redacted] response dated [redacted] to our customer questionnaire. [redacted] outlined that it considers technical capabilities, which include the range of services offered and their functionality (including the geographical availability of such services). [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁵⁴² Ofcom/ [redacted] meeting, [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁵⁴³ 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 wouldn’t 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

⁵⁴⁴ AWS website. [What is AWS](#) [accessed 20 February 2023].

Figure 5.20: Number of product categories served, based on provider websites⁵⁴⁵



Source: Ofcom analysis of provider websites and mapping of providers’ services to IaaS and PaaS categories.
 Note: IBM does not have a dedicated hybrid/multi-cloud category. Instead, related services are captured in the Integration, Networking and Containers categories.

5.223 In relation to the mid-scale providers, Oracle’s efforts to broaden its range of cloud services and features have been acknowledged by industry experts ([36]) as bringing it closer to the market leaders in terms of hyperscale cloud capabilities.⁵⁴⁶ However, as illustrated by Figure 5.20 above, Oracle offers a slightly narrower range and fewer products within most categories.⁵⁴⁷ 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).⁵⁴⁸ On the other hand, we see that Oracle offers more products in the Databases and Hybrid and Multicloud

⁵⁴⁵ AWS website. [AWS](#); Microsoft website. [Microsoft](#); Google website. [Google](#); Oracle website. [Oracle](#); IBM website. [IBM](#) [accessed 20 March 2023].

⁵⁴⁶ [36].

⁵⁴⁷ 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 provides an indication of breadth of product range.

⁵⁴⁸ Oracle website. [AI services](#) [accessed 5 March 2023]; and AWS website. [AI services](#) [accessed 5 March 2023].

categories, which likely reflects its areas of strength from its offering in on-premises business software.

- 5.224 In relation to IBM, it offers a slightly narrower range of services across IaaS and PaaS and fewer products than the hyperscalers in most categories. IBM 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.⁵⁴⁹
- 5.225 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 5.21 below.

Table 5.21: Number of cloud infrastructure products by provider – selected small-scale providers

| | Number of products |
|----------------------|--------------------|
| OVHcloud | 17 |
| Digital Ocean | 23 |
| Scaleway | 25 |

Source: Ofcom analysis of IaaS and PaaS products listed on provider websites.

- 5.226 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. There are many more ISVs compared to cloud providers, and many of these ISVs specialise in providing one type of PaaS category (e.g. data management services). For example, in 2021, most ISVs tracked by IDC were only active in one or two of IDC’s seven PaaS product categories.⁵⁵⁰ This contrasts with the hyperscalers, which in 2021 were present across all seven of IDC’s PaaS categories.
- 5.227 We observe that the hyperscalers and Oracle⁵⁵¹ 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 [X] new services per annum in net terms to their product portfolios (i.e. factoring in service discontinuations) between 2020 and 2022.⁵⁵² For example, AWS released between [X] new services in each year

⁵⁴⁹ IBM response dated 22 March 2023 to our proposed use of information dated 14 March 2023.

⁵⁵⁰ 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 2021 H2 (published April 2022).

⁵⁵¹ [X] 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. [X] 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. [X].

⁵⁵² 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 [X] 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.

between 2020-22.⁵⁵³ 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”.⁵⁵⁴

- b) [X] 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-source PaaS) to develop new features initially not available from AWS, but was made available by the time the customer developed the features.⁵⁵⁵

5.228 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. For example, 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”.⁵⁵⁶
- b) Oracle: [X] 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.⁵⁵⁷

Providers broaden their range of products and features through R&D and acquisitions

5.229 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.

⁵⁵³ 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.

⁵⁵⁴ The Stack, 2022. [119 new AWS services and features in 30 words each](#) [accessed 1 February 2023].

⁵⁵⁵ Ofcom/ [X] meeting, 1 November 2022.

⁵⁵⁶ Context Consulting research report, slides 62-63.

⁵⁵⁷ [X].

- 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

- 5.230 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 small-scale cloud providers – [X].⁵⁵⁸
- 5.231 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.^{559, 560}
- 5.232 The hyperscalers also benefit from economies of scale and scope associated with being part of a large tech conglomerate and anchor tenant benefits (see paragraphs 5.210-5.212 above).

Access to skilled labour

- 5.233 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 [X].
- 5.234 The hyperscalers may have a further advantage due to their major non-cloud digital businesses, which may further 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. 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.
- 5.235 When trying to build their technical expertise, smaller providers and entrants may find it difficult to attract highly specialised technical staff which is a scarce resource. Hyperscalers with the financial capabilities to offer competitive salaries and established reputations as leaders in their field may be more able to attract highly specialised technical profiles. A smaller cloud provider ([X]) explained that a new entrant might find it difficult to attract

⁵⁵⁸ Based on analysis of data collected via responses to s.174 notices.

⁵⁵⁹ 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, pp. 74-78

⁵⁶⁰ 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 GAFAs'. In: *Digital Dominance: The Power of Google, Amazon, Facebook, and Apple*. Oxford University Press, New York, 2018, pp. 21-49.

highly specialised technical staff given: i) staff shortages in the cloud market; and ii) the hyperscalers offering extremely comfortable hiring conditions that new entrants might not be able to provide.⁵⁶¹

5.236 Lastly, the hyperscalers have the financial capabilities to secure the best talent by acquiring other companies. In this respect, AWS made important strategic investments at a very early stage of the cloud services market. For example:

- a) In 2015 Amazon acquired chip designer Annapurna Labs for an estimated \$350-400m.⁵⁶² The “Annapurna Labs team” is responsible for building innovation in silicon and software for AWS customers. Amazon explains that this team is behind the innovations discussed above: “[b]ecause of our teams’ breadth of talent, we’ve been able to improve AWS cloud infrastructure in networking and security with products such as AWS Nitro, Enhanced Network Adapter (ENA), and Elastic Fabric Adapter (EFA), in compute with AWS Graviton and F1 EC2 Instances, in machine learning with AWS Neuron, Inferentia and Trainium ML Accelerators, and in storage with scalable NVMe.”⁵⁶³
- b) In 2017 Amazon acquired the imaging processing technology of Intelligent Imaging Solutions GmbH (“IIS”), a 4-person image processing technology company based in Tübingen, Germany.⁵⁶⁴ [REDACTED].⁵⁶⁵

Acquisitions

5.237 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.⁵⁶⁶

5.238 We have collated cloud-related acquisition data from five of the largest cloud providers (Amazon, Microsoft, Google, IBM, and Oracle) for the period 2018 – 2021. The number of acquisitions made in this period ranges from 11 ([REDACTED]) to 23 ([REDACTED]). [REDACTED] and [REDACTED] have made fewer acquisitions than [REDACTED] but more than [REDACTED] and [REDACTED]. [REDACTED]. 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.

5.239 We note that despite their different 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.

⁵⁶¹ [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], question [REDACTED].

⁵⁶² DataCenter Knowledge, January 2015, [Amazon Buys Stealthy Israeli Chip Startup Annapurna Labs](#).

⁵⁶³ Amazon website. [Annapurna Labs](#) [accessed 1 February 2023].

⁵⁶⁴ AWS response dated 9 December 2022 to our s.174 notice dated 21 October 2022, part B, question 27.

⁵⁶⁵ [REDACTED] response dated [REDACTED] to our s.174 notice dated [REDACTED], part B, question [REDACTED].

⁵⁶⁶ Based on analysis of data collected via responses to s.174 notices.

- 5.240 Examining some recent case studies illustrates 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 2022,⁵⁶⁷ which has enabled it to expand its Chronicle Security Operations software suite.⁵⁶⁸ [§].⁵⁶⁹ This is reflected in its \$2.6bn acquisition of data analytics start-up Looker in 2020.⁵⁷⁰
 - 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.⁵⁷¹
 - c) Microsoft acquired Nuance, a provider of speech recognition and AI solutions, in 2022 for \$20bn.⁵⁷² In a press release, Microsoft said that the acquisition represented the 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.⁵⁷³
 - 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.⁵⁷⁴
 - 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.⁵⁷⁵ [§].⁵⁷⁶
- 5.241 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

- 5.242 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.

⁵⁶⁷ TechCrunch, 2022. [Google closes \\$5.4B Mandiant acquisition](#) [accessed 28 February 2023]

⁵⁶⁸ TechCrunch, 2022. [Google looks to boost its security cred in the cloud](#) [accessed 7 March 2023]

⁵⁶⁹ [§] response dated [§] to the s.174 notice dated [§], question [§].

⁵⁷⁰ TechCrunch, 2020. [Google closes \\$2.6B Looker acquisition](#) [accessed 28 February 2023]

⁵⁷¹ Financier Worldwide, 2022. [Oracle agrees to acquire Cerner for \\$28bn](#) [accessed 28 February 2023]

⁵⁷² TechCrunch, 2022. [After clearing all regulatory hurdles, Microsoft closes \\$20B Nuance deal](#) [accessed 28 February 2023]

⁵⁷³ Microsoft press release, 2021. [Microsoft accelerates industry cloud strategy for healthcare with the acquisition of Nuance](#) [accessed 1 March 2023]

⁵⁷⁴ 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 28 February 2023]; Microsoft website, 2020. [Microsoft announces agreement to acquire Affirmed Networks to deliver new opportunities for a global 5G ecosystem](#)

⁵⁷⁵ Red Hat website, 2019. [IBM closes landmark acquisition of Red Hat for \\$34 billion; defines open, hybrid cloud future](#) [accessed 28 February 2023]; Softchoice, 2019. [4 takeaways from IBM's Red Hat Acquisition](#) [accessed 28 February 2023]

⁵⁷⁶ [§] response dated [§] to the s.174 notice dated [§], question [§].

- 5.243 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.⁵⁷⁷ [§<].⁵⁷⁸ [§<] 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”.⁵⁷⁹

Network effects

- 5.244 Cloud providers can add breadth and variety to their product offering by attracting ISVs to offer services on top of their cloud infrastructure. This benefits customers by allowing them to select ISV services and combine them with the cloud providers’ first-party services when building their cloud solutions.
- 5.245 Indirect network effects may exist for ISV 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 the cloud of a particular provider may increase with the number of users they can access with that provider.
- 5.246 The presence of indirect network effects can act as a barrier to entry and expansion because 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.
- 5.247 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 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.
- 5.248 In its response to our CFI, Microsoft noted that – unlike other IT settings⁵⁸⁰ – 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.⁵⁸¹ However, based on the evidence we have received we consider that network effects are in fact an important feature of the cloud market. In particular:

⁵⁷⁷ Google website. [Google Maps Platform now integrated with the GCP Console](#) [accessed 27 February 2023].

⁵⁷⁸ [§<] response dated [§<] to the s.174 notice dated [§<], question [§<].

⁵⁷⁹ [§<] response dated [§<] to the s.174 notice dated [§<], page [§<].

⁵⁸⁰ 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.

⁵⁸¹ Microsoft response to the CFI, page 12.

- 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.
- 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.⁵⁸² 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.⁵⁸³ 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, two popular ISVs ([redacted]) said that expected customer demand was one of their key considerations when deciding which cloud provider to deploy their services on.⁵⁸⁴

5.249 In addition, the evidence received from cloud providers and ISVs indicates that indirect network effects may act as a barrier to entry and favour more popular cloud providers. In particular:

- a) Popular ISVs ([redacted]) we have engaged with have only integrated their services with AWS, Microsoft or Google.⁵⁸⁵ Some of these ISVs ([redacted]) 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, [redacted] launched on AWS first, Azure two years later and on Google four years later. Similarly, [redacted] launched on AWS first and on Azure and Google one year later.⁵⁸⁶ Moreover, [redacted] stated that achieving interoperability with an additional cloud infrastructure would involve significant costs.⁵⁸⁷
- 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

⁵⁸² Microsoft response dated 18 November 2022 to the s.174 dated 21 October 2022, part A question 31.

⁵⁸³ OVHcloud response dated 17 November 2022 to s.174 notice of 27 October 2022, part A question 31b, page 32.

⁵⁸⁴ [redacted] response dated [redacted] to the s.174 dated [redacted], part A question [redacted]. [redacted] response dated [redacted] to the s.174 dated [redacted], part A question [redacted].

⁵⁸⁵ [redacted] response dated [redacted] to the s.174 dated [redacted], part A question [redacted]. [redacted] response dated [redacted] to the s.174 dated [redacted], part A question [redacted]. [redacted].

⁵⁸⁶ More specifically, [redacted] ([redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]). [redacted] ([redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]).

⁵⁸⁷ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

Rancher offered support for AWS first and then expanded their products to Microsoft, Google and others. Moreover, we understand Terraform may not offer the full set of functionalities on smaller clouds.

- c) [X] submitted that – while this may be 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.⁵⁸⁸
- d) In line with the above, [X] 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.⁵⁸⁹
- e) 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.⁵⁹⁰ For example, ([X]) are only available on the marketplaces of the hyperscalers.⁵⁹¹

5.250 This type of indirect network effects may be mitigated by customers’ ability to integrate services hosted on different clouds (i.e. ‘integrated multi-cloud’).⁵⁹² 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. For example, customers might be more willing to use a storage service of a smaller cloud provider if they could still access their preferred analytics service from one of the hyperscalers. However, the barriers to multi-cloud we have identified above suggests that such solutions cannot reduce the effects of these indirect network effects.

5.251 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:

- a) First, as set out earlier in this section, we consider cloud exhibits a high degree of skills specialisation (i.e. the skills of cloud developers / engineers differ across clouds). This

⁵⁸⁸ Ofcom/ [X] meeting, [X].

⁵⁸⁹ [X] response dated [X] to the s.174 notice dated [X], question [X].

⁵⁹⁰ 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.

⁵⁹¹ [X] response dated [X] to the s.174 notice dated [X], question [X]. [X] response dated [X] to the s.174 notice dated [X], part A question [X]. [X].

⁵⁹² We recognise that customers could adopt other multi-cloud architectures (i.e. ‘siload multi-cloud’ and ‘cloud duplication’). However, when these multi-cloud architectures are adopted, customers would typically distribute workloads across providers (i.e. siload multi-cloud) or only use the secondary cloud as a back-up for some of their workloads (i.e. cloud duplication). As demonstrated by shares of supply, the portion of workloads that are run on smaller cloud providers is low compared to hyperscalers. Therefore, the adoption of these multi-cloud architectures is unlikely to mitigate network effects since ISVs would still have a greater incentive to deploy their services on clouds where the majority of workloads is being run (i.e. hyperscalers).

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.

- b) Second, some customers have told us that 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, [redacted] 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.⁵⁹³ 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.

- 5.252 These types of direct network effect can further add to barriers to entry and expansion as once a cloud reaches a certain userbase size, it can be difficult for smaller cloud providers to challenge them effectively.

Provisional conclusion

- 5.253 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.
- 5.254 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.

Customer acquisition strategies

- 5.255 Existing customer relationships in adjacent markets, and offering substantial credits, appear to be important ways in which cloud providers attract new customers. In this subsection, we explore whether smaller cloud providers can benefit from existing

⁵⁹³ Ofcom/ [redacted] meeting, [redacted].

relationships to the same extent as the hyperscalers, and also whether they are able to match the substantial credits offered by the largest providers. We also consider whether historical choices of cloud provider create additional barriers to entry and expansion.

Importance of historical choices of cloud provider

- 5.256 Customers' historical choices and existing relationships with providers can significantly influence their future decisions on choice of provider. For example, one cloud provider ([REDACTED]) suggest that while factors such as pricing and the range of services offered are important in attracting customers to cloud providers, customers' historic relationships with cloud providers remains a "leading" factor in influencing their choice of cloud provider.⁵⁹⁴
- 5.257 In terms of historical choices of cloud provider, many customers will already have an established relationship with the leading providers, with AWS likely to have captured a material share of customers that were early adopters of cloud and Microsoft benefiting from and building on its strong position in enterprise software (discussed further below). When new use-cases and workloads emerge, customers that decide to single-source will likely continue to take services from their chosen cloud provider – because of the high switching costs and barriers to multi-cloud (discussed above).
- 5.258 However, our evidence indicates that in some cases, particularly for larger and more sophisticated customers, there can be competition for new workloads that are sufficiently separate to a customer's existing workload hosted by their primary cloud provider. Smaller providers that entered the market later can compete for such new workloads and compete to become the secondary provider for existing cloud customers (i.e. with AWS or Microsoft remaining as the primary provider). While this provides a way for other providers such as Google, IBM and Oracle to build their positions in the market, the scope to build scale as secondary providers appears limited. By way of illustration, [REDACTED] told us that [REDACTED].⁵⁹⁵ This is also supported by evidence collated from large customers, for example:
- a) [REDACTED].⁵⁹⁶
 - b) [REDACTED].⁵⁹⁷
 - c) [REDACTED].⁵⁹⁸
 - d) [REDACTED].⁵⁹⁹
 - e) [REDACTED].⁶⁰⁰
 - f) [REDACTED].⁶⁰¹

⁵⁹⁴ [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], question [REDACTED].

⁵⁹⁵ [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], question [REDACTED].

⁵⁹⁶ [REDACTED] response dated [REDACTED] to our customer questionnaire.

⁵⁹⁷ [REDACTED].

⁵⁹⁸ [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], questions [REDACTED].

⁵⁹⁹ [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], question [REDACTED].

⁶⁰⁰ [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], questions [REDACTED].

⁶⁰¹ [REDACTED] response dated [REDACTED] to the s.174 notice dated [REDACTED], questions [REDACTED].

Importance of existing relationships in adjacent software markets

- 5.259 Our evidence also indicates that existing relationships with providers for non-cloud services can 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).⁶⁰² Of all factors that were chosen as 'most important' existing relationships received the third highest 'votes'.⁶⁰³
- 5.260 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. For example, as mentioned in Section 4, 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 note 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.⁶⁰⁴
- 5.261 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.⁶⁰⁵ Some customers identify Azure as a natural choice when already using Microsoft for other services.
- 5.262 From the 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".⁶⁰⁶ 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.⁶⁰⁷
- 5.263 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

⁶⁰² Context Consulting research report, slide 69.

⁶⁰³ After service quality (ranked first) and value for money (ranked second). Context Consulting research report, slide 69.

⁶⁰⁴ CMA, 12 October 2022. [Microsoft/Activision phase 1 decision](#), paragraph 260.

⁶⁰⁵ Context Consulting research data tables, Q25.

⁶⁰⁶ Context Consulting research report, slide 57.

⁶⁰⁷ Context Consulting research report, slide 56.

support it provided.⁶⁰⁸ 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 [redacted].⁶⁰⁹

- 5.264 Other customers suggest there may be licensing advantages associated with using Microsoft Azure when choosing their cloud provider. For example, [redacted] 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", but it does not further explain the licensing reasons.⁶¹⁰ These licensing reasons/advantages could be in relation to lower costs, including pricing incentives. We are aware that enterprises have the ability to pay a reduced rate for Azure, if they have an existing on-premises Windows Server/SQL licence.⁶¹¹ Therefore, 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, [redacted] recognise that existing relationships may provide a degree of familiarity to customers in the form of established support and known licensing models.⁶¹²
- 5.265 Mid-scale cloud providers, such as IBM and Oracle are also known to provide on-premises services (e.g. Oracle databases), whilst the ISV VMware is known for its private cloud platform. Therefore, these cloud providers 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.⁶¹³
 - b) In our engagement with large customers, [redacted] explained that it is currently considering the use of Oracle Cloud for its (on premise) Oracle workloads due to its existing relationship with Oracle for non-cloud services.⁶¹⁴
 - 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.⁶¹⁵
 - 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".⁶¹⁶

⁶⁰⁸ Ofcom / ASOS meeting, 29 November 2022.

⁶⁰⁹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶¹⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶¹¹ Microsoft website. [Azure Hybrid Benefit](#) [accessed 2 February 2023].

⁶¹² [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶¹³ Context Consulting research data tables, Q25.

⁶¹⁴ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶¹⁵ Oracle response dated 20 March 2023 to our proposed use of information dated 14 March 2023.

⁶¹⁶ IBM response dated 6 December 2022 to the s.174 notice dated 25 October 2022, part A, question 7.

- 5.266 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.
- 5.267 There are also some suggestions that AWS and Google may benefit from existing relationships. One cloud provider ([redacted]) claims 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.⁶¹⁷ 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.⁶¹⁸ Another cloud provider ([redacted]) argues that Google leverages its existing relationships with Chief Marketing Officers,⁶¹⁹ and continue to "bundle cloud offers with Advertising and Google Workspace" to attract new customers.⁶²⁰ [redacted].⁶²¹ 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.⁶²² Furthermore, the market research finds Google and AWS benefit from relationships in other services to a lesser degree than Microsoft.⁶²³

Cloud credits

- 5.268 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.
- 5.269 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.⁶²⁴
- 5.270 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-

⁶¹⁷ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶¹⁸ The Information, 2021. [WarnerMedia Extended AWS Deal to Win Key HBO Max Concession](#) [accessed 9 March 2023].

⁶¹⁹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶²⁰ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶²¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶²² Context Consulting research report slide 70; and Context Consulting research data tables Q25.

⁶²³ 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.

⁶²⁴ Azure website. [Azure free account](#) [accessed 17 February 2023]; and IBM website. [IBM Cloud free tier](#) [accessed 17 February 2023].

ups/scale-ups and offer exclusive benefits to the venture capital firm's portfolio of companies. There is also much more variability across cloud providers. For example:

- 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.⁶²⁵
- 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.⁶²⁶ 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.⁶²⁷
- c) Oracle's website does not currently specify the monetary value of cloud credits or discounts it offers to start-ups. There is some evidence to suggest that Oracle did offer \$500 in cloud credits and 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.⁶²⁸
- d) OVHcloud's start-up programme offers up to 100k euros (or equivalent local currency) for the first year.⁶²⁹
- e) Scaleway's start-up programme offers up to 36k euros for the first year.⁶³⁰

5.271 Some smaller cloud providers ([redacted]) told us that they are unable to match the credits offered by the hyperscalers.⁶³¹ [redacted] explained that its credits are much more limited in comparison to the hyperscalers in monetary terms [redacted] and also in terms of the total number of companies it can offer credits to [redacted].⁶³² [redacted] explained it cannot afford/absorb the losses incurred during these credit periods.⁶³³

5.272 The hyperscalers issued substantial credits to UK customers in 2021: AWS issued \$[redacted] worth of credits (approximately [redacted]% of its public cloud revenue); Google issued \$[redacted] worth of credits (approximately [redacted]% of its public cloud revenue).⁶³⁴ Whilst Microsoft was unable to provide data on cloud credits issued to UK/Global customers, it did provide an

⁶²⁵ AWS website. [AWS activate](#) [accessed 17 February 2023]; Azure website. [Unlocking Azure credits as your start-up grows](#) [accessed 17 February 2023]; and Google website. [Google for start-ups cloud program](#) [accessed 17 February 2023].

⁶²⁶ IBM website. [The start-up with IBM program](#) [accessed 17 February 2023].

⁶²⁷ Aurelia Ventures website. [IBM Cloud credits](#) [accessed 6 March 2023]; and DSW Ventures website. [IBM](#) [accessed 6 March 2023].

⁶²⁸ 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, but this may be an archived document. Oracle website. [Oracle for Startups](#) [accessed 6 March 2023]

⁶²⁹ OVHcloud website. [Startup programme FAQs](#) [accessed 6 March 2023].

⁶³⁰ Scaleway website. [Why choose Scaleway?](#) [accessed 17 February 2023].

⁶³¹ Ofcom/[redacted] meeting, [redacted] and Ofcom/[redacted] meeting [redacted].

⁶³² Ofcom/[redacted] meeting, [redacted].

⁶³³ Ofcom/[redacted] meeting, [redacted].

⁶³⁴ Note that the total amount of credits issued may not have been redeemed by customers during that period. AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, part B question 8; Google response dated 22 December 2022 to the s.174 notice dated 26 October 2022, part B question 8.

estimate of cloud credits redeemed by UK customers of \$[X] (although this is expected to be lower than actual credits redeemed due to it not having a UK split for some credit programmes).⁶³⁵ These figures are high relative to other cloud providers in absolute terms and relative to their cloud revenue: IBM issued \$[X] worth of credits (approximately [X]% of its public cloud revenue) and Oracle issued \$[X] worth of credits (approximately [X]% of its public cloud revenue).⁶³⁶ Furthermore, AWS's credit issuance exceeds the UK public cloud revenue of some smaller providers (e.g. [X] UK public cloud revenue of \$[X]).⁶³⁷

- 5.273 Smaller providers may not be able to match the cloud credits offered to start-ups/scale-ups because offering credits to start-ups/scale-ups entails a certain level of risk. There is a risk that the 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 by means of scale as a cloud provider that attracts large numbers of start-up/scale-up customers may be able to diversify risk more effectively by acquiring customers across a variety of different industries.
- 5.274 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. Furthermore, customers' familiarity with the hyperscalers and a perception that they represent the safe option in cloud may also be relevant factors.⁶³⁸ Indeed, the market research found that few customers had considered using providers outside of the hyperscalers, with some customers lacking any awareness of smaller providers.⁶³⁹ 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.
- 5.275 Based on the information gathered so far, we think that the credits offered by cloud providers may not automatically be redeemable on third-party services of ISVs.⁶⁴⁰ This could create barriers to entry for ISVs who may not be able to offer competitive cloud credits.⁶⁴¹ We have not received evidence from ISVs explicitly suggesting that they are not able to match the hyperscalers' cloud credit offerings. However, we have only received

⁶³⁵ Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, part B question 8.

⁶³⁶ IBM response dated 23 December 2022 to our s.174 notice dated 25 October 2022, part B question 8; and Oracle response dated 16 December 2022 to the s.174 notice dated 31 October 2022, part B question 8.

⁶³⁷ AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, part B question 8; [X] response dated [X] to the s.174 notice [X], part B question [X].

⁶³⁸ Context Consulting research report, slides 64-65.

⁶³⁹ Context Consulting research report, slides 64-65.

⁶⁴⁰ For example, in their community blog MongoDB advise a customer asking about use of AWS credits to send a separate email to MongoDB start-up team, which suggest that use of AWS credits on MongoDB is not automatic. MongoDB website. [Developer community](#) [accessed 16 March 2023]. Moreover, this article from Google explaining their partnership with MongoDB suggests that when credits partnerships are created, they are explicitly marketed. Google website. [Google cloud blog](#) [accessed 16 March 2023].

⁶⁴¹ We understand that popular ISVs such as MongoDB and Snowflake offer their own credits schemes to customers. For example, MongoDB has a start-up programme for start-ups that have been referred by venture capitalist or accelerator partner firms. MongoDB website. [Start-ups](#) [accessed 16 March 2023]. Snowflake offers free trials. Snowflake website. [Snowflake free usage](#) [accessed 16 March 2023].

information from a relatively small number of ISVs to date, and therefore seek further evidence from a wider range of ISVs on this issue.

- 5.276 We also have some evidence of cloud credits being used to lower the cost of migration from on-premises: [redacted] and [redacted] told us that they each received substantial cloud credits of \$[redacted]m and \$[redacted]m respectively, over a 5-year term from Google to be used on professional services to assist with services such as data migrations.⁶⁴² [redacted] told us that Oracle offered it cloud credits as an incentive to migrate workloads from Oracle on-premises solutions to Oracle's cloud.⁶⁴³ 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 and/or additional workloads for a secondary cloud provider.

Provisional conclusion

- 5.277 Overall, our evidence suggests that the hyperscalers (Microsoft and AWS in particular) are better able to build the scale of their customer bases relative to smaller providers.
- 5.278 Many customers already have an established relationship with AWS or Microsoft in cloud and given there are 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 new workloads that are sufficiently separate to a customer's existing cloud usage. But the potential to gain scale in this way may be limited. 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 is likely to have contributed to their ability to enter the market and gain a small market share.
- 5.279 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. Smaller providers may not be able to match the cloud credits offered to start-ups/scale-ups because they have lower volumes of new start-up customers which may reduce the ability to diversify risk. In addition, migration credits are used to attract customers with established businesses, and may be a less risky way for smaller providers to win new customers.

Provisional conclusion on barriers to entry and expansion

- 5.280 We provisionally conclude that the factors we identify in this subsection combine to pose material barriers to entry and expansion in the provision of cloud infrastructure services. In particular in relation to broad ecosystems of services powered by a global network of data centres.

⁶⁴²[redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted] question [redacted].

⁶⁴³ [redacted] response dated [redacted] to the s.174 notice dated [redacted] question [redacted].

- 5.281 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:
- 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.
 - 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.
- 5.282 Other providers appear to have faced more barriers to grow scale in cloud infrastructure services, reducing the effective constraint that they exert on the market leaders:
- 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, switching and multi-clouding costs 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.
 - 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, 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, e.g. compared to the hyperscalers, there will be a smaller set of ISV services available on their clouds and a smaller pool of engineers trained to work with these providers. This suggests they pose only a limited constraint on AWS and Microsoft when customers choose their cloud ecosystem provider.
 - 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.

5.283 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. However, the potential reliance of ISVs on the hyperscalers may raise other potential issues, which we discuss in the next subsection).

Question 5.2: Do you agree with our analysis of potential barriers to entry and expansion?

Hyperscalers' relationship with ISVs

Hyperscalers' relations with ISVs could impact competition

- 5.284 The hyperscalers can play a number of different roles with ISVs:
- a) Hyperscalers can act as **suppliers to ISVs**. ISVs may rely on the hyperscalers' cloud infrastructure services to run cloud services that may complement or compete with those offered by the hyperscalers.
 - b) Hyperscalers can also act as **distributors of ISV services**. Hyperscalers can provide ISVs with a route to market. This could be through directly selling ISV services, offering ISVs a platform through which to sell their services (such as a marketplace), or by offering ISVs access to customers.
 - c) Hyperscalers operate and sell services across all layers of the cloud stack. As such, the hyperscalers can also **compete directly with ISVs**, offering services at the same layer of the cloud stack.
- 5.285 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 hosted on their clouds. This may create a potential conflict of interest and give the hyperscalers the opportunity to provide their own services with a competitive advantage over ISVs' competing services, ultimately increasing barriers to entry and expansion for ISVs.
- 5.286 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, the hyperscalers may deny, restrict, or increase costs of access to the cloud infrastructure services ISVs need to effectively run their services in the cloud. This could be done through technical mechanisms (such as not providing the necessary public APIs), and / or commercial mechanisms (such as raising prices for ISV use of the hyperscalers cloud infrastructure services). As a result, ISVs might see a reduction in the discoverability, quality and / or ease of use of their services, compared to equivalent hyperscalers services not subject to the same limitations.

b) In their role as distributor (particularly via marketplaces), the hyperscalers might be able to raise ISVs costs (for example, by increasing marketplace commission fees), self-preference their own services over ISVs' (such as through making their own services more prominent), or gain information advantages (such as gaining an understanding of which products perform well).

5.287 In the remainder of this subsection, we separately consider the hyperscalers' roles as input suppliers for, and distributor of, ISVs services. For each of those roles we assess: (i) the extent to which ISVs are reliant on the hyperscalers; and (ii) the extent to which the 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 bases

5.288 Evidence from ISVs indicate that they typically deploy their services across the three hyperscalers, at least initially, to access their respective user bases. For example, as discussed above, two popular ISVs ([redacted]), submitted that expected customer demand was one of the key considerations when deciding which cloud to deploy their services on.⁶⁴⁴ These ISVs also noted that they deployed their services on AWS first and expanded to Google and Microsoft one to four years later.

5.289 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 above. [redacted] 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.⁶⁴⁵ Similarly, [redacted] 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 the hyperscaler regions due to potential latency issues which would be critical for use-cases requiring a [redacted].⁶⁴⁶

5.290 In principle, 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.

ISVs can access the hyperscalers' cloud infrastructure services, but we have heard some concerns that the hyperscalers may favour their own services

5.291 The hyperscalers submitted that they provide the technical information and support needed by any third parties (including ISVs) to fully interoperate with their cloud

⁶⁴⁴ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]. [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

⁶⁴⁵ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

⁶⁴⁶ [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted].

infrastructure services or hardware. This includes, in particular, documentation about the services available and information on how third parties can interoperate with them (e.g. public APIs, and protocols). More specifically:

- a) Google said that in order to facilitate interoperability with GCP, it offers Google Cloud APIs (i.e. programmatic interfaces to Google cloud products and 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 (e.g. via server apps, mobile apps).⁶⁴⁷
- b) AWS said that it makes many of its SDKs and APIs publicly available under open-source licenses, so that customers and third parties may freely use, modify and distribute them without restriction. 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.⁶⁴⁸
- c) 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.⁶⁴⁹

5.292 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) 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 cloud providers.⁶⁵⁰
- b) 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.⁶⁵¹ 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.⁶⁵²

⁶⁴⁷ Google response dated 23 November 2022 to the s.174 notice dated 26 October 2022, part A question 32a, page 59-60.

⁶⁴⁸ 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.

⁶⁴⁹ Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, part A question 31.

⁶⁵⁰ Google response dated 23 November 2022 to the s.174 notice dated 26 October 2022, part A question 32.

⁶⁵¹ AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, part A question 31b, paragraph 31.11.

⁶⁵² AWS response dated 21 November 2022 to the s.174 notice dated 24 October 2022, part A question 31a(iii).

- c) Microsoft said that it makes the same features of Azure available to 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. It is critical to Microsoft's ability to compete to make Azure an attractive platform on which third-party services providers can deploy their solutions.⁶⁵³

5.293 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) [X] 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.⁶⁵⁴
- b) [X] submitted that all the hyperscalers open their APIs sufficiently to allow it to integrate its software ([X]) and achieve similar performance to the hyperscalers first-party services in relation to core software functionality.⁶⁵⁵
- c) [X] 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.⁶⁵⁶

5.294 However, [X] have given examples of cases where AWS and Microsoft only allow access to their proprietary APIs for the benefit of their first-party products, thereby negatively affecting discoverability and ease of use for [X] services.⁶⁵⁷ In particular, [X] explained that:

- a) [X] services appear less discoverable than their first-party counterparts. For example, AWS first-party counterparts feature prominently in the AWS console⁶⁵⁸ whereas [X] services need to be discovered and installed from a crowded marketplace. Hence, customers may be less likely to take-up [X] service.
- b) Customers are typically required to complete additional steps to set-up and manage these services compared to a case where they want to add a first-party service. For example, since ISVs cannot connect into all internal APIs, customers may be required to perform additional steps to manage authentication, authorisation and network connectivity. This may increase friction and discourage take-up of ISVs.
- c) Their services would appear less embedded into AWS and Microsoft clouds which may affect the overall user experience. For example, unlike AWS's first-party [X] services ([X]) where the underlying infrastructure with all its complexities is managed by AWS

⁶⁵³ Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, part A question 31a.

⁶⁵⁴ [X] response dated [X] to the s.174 notice dated [X], part A question [X].

⁶⁵⁵ Ofcom/[X] meeting, [X].

⁶⁵⁶ Ofcom/[X] meeting, [X].

⁶⁵⁷ Ofcom/[X] meeting, [X]. Ofcom/[X] meeting dated [X].

⁶⁵⁸ AWS Console is a web application that allows users to access Amazon Web Services. The console can be considered the backbone or basic web infrastructure through which Amazon Web Services can be accessed. Source: [Javapoint](#) [accessed 16 March 2023].

and hidden away from the user, a customer has much more visibility into how [X] services are deployed which makes them appear less polished. The current APIs do not allow [X] to achieve the same level of integration which puts them at a disadvantage compared to native first-party services.

5.295 In addition, two ISVs ([X]) said that the hyperscalers may also be using other technical mechanisms to exploit their position as full-stack 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 tweak them such that customers wanting to switch need to rewrite much of their code. Customers may be induced to take-up these first-party adaptations (instead of the original open-source) because – as noted above – AWS and Microsoft make them more prominent, integrated, easy to use and / or because they may not fully disclose these adaptations such that some customers may not be aware these are different from the original open-source version. More specifically:

- a) [X] explained that AWS and Microsoft have developed proprietary [X] services which they sell as compatible with the [X] 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 [X] service. However, according to [X], the compatibility of these first-party services is limited, meaning that a customer using such [X] compatible services would need to considerably reconfigure their application if they wanted to switch to [X].⁶⁵⁹
- b) Similarly, [X] explained that the hyperscalers and other large platform vendors unfairly compete with best-of-breed innovators within their ecosystems by building their first-party services on open standards and open APIs but tweaking them. 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 much of their code.⁶⁶⁰

5.296 We have also received evidence that the hyperscalers may have an advantage as full-stack providers by offering bundles of discounted services that ISVs are not able to match:

- a) [X] told us that the main expansion challenge it has faced is the aggressive discounts the hyperscalers offer on comparable first-party services.⁶⁶¹ [X] explained that the hyperscalers can offset an aggressive discount on one product with the additional revenue generated by other products in their ecosystems. However, this strategy is unfeasible for [X] since it offers a single product.
- b) [X] said that adoption of its service has been limited because procuring multiple services from a cloud provider can be cheaper. This is because when customers source all their services from a single cloud provider they would only pay a single fee to that cloud provider. Whereas if customers decided to use [X] services in combination with

⁶⁵⁹ Ofcom/[X] meeting dated [X].

⁶⁶⁰ Ofcom/[X] meeting, [X].

⁶⁶¹ [X].

those of a cloud provider they would also need to pay the fee charged by [redacted] to use its service.

- 5.297 Lastly, two ISVs ([redacted]) indicated that they are generally wary of potential future risks of the hyperscalers self-preferencing their own services, including through technical restrictions. However, they have not presented any evidence of these risks currently materialising.⁶⁶²

Provisional conclusion

- 5.298 The evidence we have received so far suggests that ISVs materially rely on the hyperscalers' cloud infrastructure to access customers on their respective ecosystems. The feedback received from ISVs suggests that the hyperscalers are currently providing the technical information and support needed by ISVs to run their services in the hyperscalers' clouds. However, we have also identified how the hyperscalers might be able to discourage the take-up of rival ISVs services through technical restrictions,⁶⁶³ and we have seen some evidence of specific cases where AWS and Microsoft may be doing this.
- 5.299 Our provisional view is that the hyperscalers still appear to have strong incentives to attract ISVs into their ecosystems 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 ISVs' services in their ecosystems.
- 5.300 However, we have only received information from a relatively small number of ISVs to date, and therefore seek further evidence from a wider range of ISVs on this issue. We also acknowledge that the risk of the hyperscalers restricting the availability or quality of ISVs' services in their ecosystems may increase if the incentives for the hyperscalers to draw ISVs onto their ecosystem were to decrease in future.

Hyperscalers as distributors of ISV services

- 5.301 The hyperscalers can distribute ISV 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 that infrastructure).
- 5.302 The hyperscalers told us that it is rare for them to directly sell the services of ISVs to business customers.⁶⁶⁴ However, as explained above, the hyperscalers offer co-sell support to ISVs, and each hyperscaler offers a marketplace (an online platform), where providers

⁶⁶² [redacted] response dated [redacted] to the s.174 notice dated [redacted], part A question [redacted]. [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], part A question [redacted]. [redacted].

⁶⁶³ These may include hyperscalers technically restricting ISVs' access to existing or new proprietary features or services.

⁶⁶⁴ [redacted]. See [redacted] response dated [redacted] to s.174 notice dated [redacted], part A question [redacted]; [redacted] response dated [redacted] to s.174 notice dated [redacted], part A question [redacted]; [redacted] response dated [redacted] to s.174 notice dated [redacted], part A question [redacted].

can offer services to customers which run on the underlying infrastructure of the marketplace provider.⁶⁶⁵

- 5.303 Marketplaces can allow ISVs to access a broader range of customers, and so can support their ability to compete against 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 cloud 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 expansion for ISVs and weaken the ability of ISVs to compete with the hyperscalers for PaaS and SaaS services.
- 5.304 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.
- 5.305 Whether ISVs rely on the hyperscalers as a route to market will depend on the extent to which customers view different discovery and purchase routes as substitutes, and the extent to which the 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, if there were groups of customers which only used a single route to market, and if these groups of customers were sufficiently important for ISVs, then ISVs may be reliant on the hyperscalers as a route to market.

At present, ISVs do not significantly rely on marketplaces to access customers

- 5.306 At present, only a small proportion of customers purchase services through marketplaces. The market research found that only 13% of IaaS/PaaS users in the market research purchase third-party services through marketplaces.⁶⁶⁶ 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.⁶⁶⁷ This implies that most customers are purchasing ISV services through other channels.
- 5.307 We also estimate that only a small proportion of ISV sales are made through marketplaces. We estimate that less than 10% of total ISV sales (PaaS and SaaS) in the UK are transacted via a marketplace.⁶⁶⁸ Even for those ISVs which we understand make a higher proportion of

⁶⁶⁵ We understand that a number of other smaller cloud providers 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 ISV services, we do not consider smaller cloud provider marketplaces further.

⁶⁶⁶ Analysis of Context Consulting research data, Q45 and Q46. We do not know what proportion of their ISV purchases customers make through marketplaces.

⁶⁶⁷ Analysis of Context Consulting research data, Q45 and Q46.

⁶⁶⁸ We calculated two estimates of total PaaS + 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.

their sales through marketplaces, marketplace sales still make up only a minority of their sales. This implies that ISVs can and do sell their services via other channels. [redacted] notes that “many sellers don’t yet view marketplaces as a strategic new customer acquisition channel”.⁶⁶⁹

- 5.308 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 do not have strong evidence on the importance of marketplaces for smaller ISVs, and we are keen to hear from any ISVs on this point.
- 5.309 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.⁶⁷⁰ 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.⁶⁷¹ 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 the ISV.⁶⁷² But others have told us that they do not use marketplaces and that they prioritise direct procurement relationships with vendors.⁶⁷³ Other customers make very limited use of marketplaces, or do not use them at all.⁶⁷⁴ Many of these customers prefer to procure directly with vendors in order to control cost. Even those customers that currently make use of marketplaces tell us that they are able to use multiple routes to market, going direct to vendors if they want to.⁶⁷⁵ We have not seen any specific evidence that any customer groups are particularly reliant on marketplaces as a purchase channel.⁶⁷⁶
- 5.310 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

⁶⁶⁹ See [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted].

⁶⁷⁰ Including [redacted]. See [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; [redacted] response dated [redacted] to our customer questionnaire, question [redacted]. 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 10 February 2023].

⁶⁷¹ See [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; [redacted] response dated [redacted] to our customer questionnaire, question [redacted]; and [redacted] response dated [redacted] to our customer questionnaire, question [redacted]. “Billing for existing services” was the largest single purpose for using marketplaces (64% of marketplace users). Context Consulting research data tables, Q46.

⁶⁷² [AWS Marketplace customers use private offers, big customers - Protocol](#) [accessed 9 March 2023].

⁶⁷³ For example, [redacted] told us that it would only choose to purchase products through a cloud marketplace if a direct relationship is not possible. [redacted] response dated [redacted] to our customer questionnaire, question [redacted].

⁶⁷⁴ This includes [redacted]. See [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶⁷⁵ For example, [redacted] told us that it is normally able to purchase directly from the vendor. See [redacted] response dated [redacted] to our customer questionnaire.

⁶⁷⁶ [redacted] told us that it has not observed a particular pattern of types of customers procuring through particular distribution channels. [redacted] response dated [redacted] to s.174 response dated [redacted], question [redacted].

whether marketplaces have an important role in discovery is mixed. The market research suggests that more IaaS/PaaS users use marketplaces to discover new services (25%) than to purchase ISV services (13%).⁶⁷⁷ 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.⁶⁷⁸ This suggests that marketplaces could have a greater role in discovering service options, which may then be purchased either through the marketplace or elsewhere.

- 5.311 Customers differ in whether they use marketplaces for discovering new services. [redacted] says it uses marketplaces to discover services but may then buy direct, but in contrast, [redacted] state that they use marketplaces to purchase services they are already aware of.⁶⁷⁹ Further, AWS argues that marketplaces are rarely used for casual shopping, and Microsoft argues that sophisticated customers are using the marketplace to fulfil specific cloud services needs with solutions the customers have already determined before visiting the marketplace.⁶⁸⁰
- 5.312 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.⁶⁸¹ Overall therefore, 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 owners

- 5.313 Given our provisional finding that, in general, ISVs are not currently reliant on marketplaces to access customers, marketplace owners 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.
- 5.314 Marketplace owners do not require ISVs to exclusively list on their marketplace.⁶⁸² Marketplace owners' mechanisms for organising and displaying solutions to customers do

⁶⁷⁷ Analysis of Context Consulting research data, Q45 and Q46.

⁶⁷⁸ Analysis of Context Consulting research data, Q46.

⁶⁷⁹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶⁸⁰ 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.

⁶⁸¹ Analysis of Context Consulting research data, Q46. The market research does not tell us how many customers discover through marketplace but then buy direct.

⁶⁸² 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

not appear to preference their own services – instead basing this on search experience, customer popularity or benefits programs.⁶⁸³ We have not heard any specific concerns from ISVs about the prominence or ranking of their services on marketplaces. AWS claims that it is transparent and neutral with accurate and unbiased product information,⁶⁸⁴ and we note that almost all ([redacted]%) of its marketplace sales are third-party – which suggests it does not self-preference.⁶⁸⁵ Google stated that it [redacted].⁶⁸⁶ 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 draw down through marketplace sales is limited (e.g. [redacted]), or is less than the draw down for purchases of first-party (e.g. [redacted]).⁶⁸⁷

- 5.315 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.⁶⁸⁸ 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 restrict service product and engineering teams from accessing transactional data from its marketplace.⁶⁸⁹ We welcome further evidence on this.
- 5.316 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

Marketplace. However, this does not explicitly prevent vendors from selling their services through other routes. [SaaS product guidelines - AWS Marketplace \(amazon.com\)](#) [accessed 14 March 2023]. [redacted] told us that listing on a particular marketplace does not prevent it from selling outside of such marketplace. See [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶⁸³ See [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]. Microsoft organises some apps as “Featured Apps” based on [redacted] and badges certain offerings as “preferred solutions” if they meet certain criteria. See [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], and [Learn about the Microsoft preferred solution badge | Microsoft Learn](#) and [Microsoft Azure Marketplace](#) [accessed 1 March 2023].

⁶⁸⁴ See AWS response dated 28 November 2022 to the s.174 notice dated 24 October 2022, Annex Q33.1.15, page 8.

⁶⁸⁵ See AWS response dated 9 December 2022 to s.174 request dated 24 October 2022.

⁶⁸⁶ See Google response dated 23 November 2022 to s.174 request dated 26 October 2022, annex 33.22, p.1.

⁶⁸⁷ See [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted], [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]. We also note that [redacted] has raised concerns that differentiated treatment in relation to marketplace spend commitments impacts on its ability to compete (Ofcom/ [redacted] meeting, [redacted]).

⁶⁸⁸ [redacted] 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 [redacted] was able to tell that some of its customers [redacted]. Ofcom/ [redacted] meeting, [redacted]. 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](#), page 63 [accessed 29 September 2022].

⁶⁸⁹ [redacted]. See [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted]; and [redacted] response dated [redacted] to our follow-up email dated [redacted] concerning the s.174 notice dated [redacted], question [redacted].

3%⁶⁹⁰ and we understand that [redacted], the average listing fee AWS receives [redacted],⁶⁹¹ and can be low-single digits for some ISVs.⁶⁹² [redacted].⁶⁹³ 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 ISV 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.⁶⁹⁴ Additionally, [redacted]⁶⁹⁵ and [redacted].⁶⁹⁶ Evidence we have gathered suggests that marketplace owners are considering and responding to changes made by other hyperscalers to their marketplaces and monitoring rivals marketplaces.⁶⁹⁷

Provisional conclusion

- 5.317 The evidence received so far suggests that, at present, ISVs 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 or behaviour within their marketplaces which could act as a barrier to entry or expansion for ISVs.
- 5.318 That said, the hyperscalers appear focused on increasing uptake of their marketplaces by their customers and ISVs (see Section 4, paragraph 4.69), and we anticipate that marketplaces are likely to grow in importance as a route to market. We are aware that the balance of the hyperscalers' incentives may change going forward and the risk of foreclosure via marketplaces may increase. We consider this risk further in Section 6.
- 5.319 At this stage we have received input from a small number of mainly larger ISVs which may not reflect the experience of other ISVs that may be in a different position. We will continue to gather evidence and views on the relationship between the hyperscalers and ISVs, and would particularly welcome further input from stakeholders including other ISVs on these topics.

Question 5.3: Do you agree with our analysis of the hyperscalers' relationship with ISVs? As part of this, please provide your views on whether our analysis of the hyperscalers relationship with ISVs applies to both larger and smaller ISVs.

⁶⁹⁰ [Partners Cheer 'Substantial' Microsoft Marketplace Fee Cuts | CRN](#) and [Google Cloud doubles-down on ecosystem in 2022 to meet customer demand | Google Cloud Blog](#). [redacted]. See [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶⁹¹ [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted]; and Ofcom analysis of [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶⁹² [AWS Marketplace customers use private offers, big customers - Protocol](#) [accessed 9 March 2023].

⁶⁹³ For example, [redacted] has negotiated fees of [redacted]%. [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁶⁹⁴ [AWS Marketplace customers use private offers, big customers - Protocol](#) [accessed 9 March 2023]. [redacted] ([redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted]).

⁶⁹⁵ [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted].

⁶⁹⁶ [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted].

⁶⁹⁷ For example, see [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], [redacted].

6. Potential competition concerns

- 6.1 In this market study, our aim is to assess whether any feature of the market, or the behaviour of providers, could dampen competition between suppliers and therefore have an adverse effect on competition leading to higher prices, lower quality products or less innovation.
- 6.2 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.
- 6.3 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.
- 6.4 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 innovation can directly affect the quality and range of choice UK consumers have of online services powered by the cloud.
- 6.5 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 (Section 5). Below, we set out our provisional views on how well competition is working in cloud infrastructure at present based on our findings in Sections 4 and 5. Second, we 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

- 6.6 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.

- 6.7 Due to the difficulties of switching and multi-cloud once a customer is established with a cloud provider, competition in cloud infrastructure is currently centered 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 later earn higher revenues as the credits and free trial periods come to an end and acquired customers begin to take up further services.
- 6.8 Our evidence indicates that AWS 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 base. This is reflected in the market shares we have estimated, that show that AWS and Microsoft accounted for [X] [60% to 70%] of UK revenues across IaaS and PaaS in 2021. Google is their closest competitor and has been able to grow its presence in the UK offering a broad range of cloud infrastructure services that rival that of the market leaders. However, it had a significantly lower share of revenues in 2021 - around [X] [5% to 10%]. There are also several smaller cloud providers present, including large technology companies such as Oracle and IBM, who both had around [X] [0% to 5%] share in 2021. These providers are more distant competitors, because of the difficulty of building a rival ecosystem of products delivered over a global network of data centres.

There is limited scope for workload competition amongst cloud providers

- 6.9 While a material fraction of our market research respondents reported using multiple providers (multi-cloud), this is often where different departments procure cloud independently, or where businesses use different clouds following an acquisition. In such cases, there may not be much active competition between a company's different providers for their existing workloads.
- 6.10 However, there appears to be some scope for competition 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 second provider to serve a new use-case. This can benefit such customers by giving them access to 'best-in-breed' services and may strengthen their bargaining position when renewing their contract with their main provider. Our evidence suggests that due to barriers to integrated multi-cloud (such as egress fees), such competition is likely limited to 'siloes' use-cases. This is where applications/data do not have to be integrated/transferred with those of the primary provider.
- 6.11 In addition, there can be a degree of workload competition within a cloud provider's ecosystem, as customers can choose between the first-party services of cloud providers and third-party services of ISVs to build their workloads. In PaaS, we estimate that ISVs collectively account for [X] [30% to 40%] of revenues and this is largely made up of a long tail of companies that each have a small share and tend to offer products in one or a small number of product categories.

While customers benefit from cloud infrastructure services, some market outcomes suggest competition is not working well

- 6.12 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:** Providers are investing to innovate and build quality products to attract new customers. Our evidence indicates that product quality, number of features and the range of products are key parameters customers consider when choosing a provider. We have found that the hyperscalers are continually developing new products to add to their existing portfolios and are at the forefront of cutting-edge technologies such as ML and AI. We also see evidence of the hyperscalers developing the features of existing products and investing in making physical infrastructure components (e.g. chips) more powerful and efficient.
 - b) **Customer choice of ISVs:** Customers have a broad range of ISV products on each of the hyperscaler's infrastructure 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 the hyperscalers benefit from 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 their own container services and container orchestration platforms, including some based on open-source products, e.g. Kubernetes. Other examples include hyperscalers offering some database services based on open-source database technologies such as PostgreSQL, MariaDB/MySQL and Redis. 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 (as discussed in more detail below).
 - d) **Pricing trends for compute:** The industry literature reports that list prices of IaaS services (such as compute) have generally decreased over time, reflecting improvements in performance and efficiency of underlying hardware.⁶⁹⁸ For example, it is estimated that general purpose AWS virtual instances have decreased in price by 50% since their first generation.⁶⁹⁹
- 6.13 However, despite providers competing to win new customers, 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:

⁶⁹⁸ IaaS Pricing Patterns and Trends, 17 December 2021. <https://redmonk.com/rstephens/2021/12/17/iaas-pricing-2021/> [accessed 2 March 2023].

⁶⁹⁹ E.g. Cloud generations drive down prices (2022) <https://journal.uptimeinstitute.com/cloud-generations-drive-down-prices/>, [accessed 2 March 2023].

- 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).⁷⁰⁰ While our market research found that switching is less cumbersome for some, we found that only c.20% of customers have switched providers⁷⁰¹ – which is likely to include switching within a cloud, 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.
- b) **Integrated multi-cloud is uncommon:** at least half of IaaS/PaaS users in our market research only use a single cloud provider.⁷⁰² Of the customers that use multiple providers, our engagement with large customers and intermediaries indicated that they typically run distinct workloads on separate clouds and that there are few occasions where cloud solutions are built by integrating clouds with data being transferred between them. This indicates that some customers are limited to using a single provider for tightly integrated use-case which restricts their choice of services and may limit their ability to build their preferred solutions.
- c) **AWS and Microsoft have a persistently high share of supply:** we estimate that between 2019 and 2021, AWS and Microsoft’s combined share of UK IaaS and PaaS revenues was stable, and in 2021 stood at [3<] [60% to 70%], with AWS’s share falling marginally and Microsoft’s growing slightly over this period. Although we have seen Google grow steadily, its share remains considerably smaller than the market leaders. Beyond the hyperscalers, the next largest providers have considerably smaller shares and overall they are losing share.
- d) **Evidence of high profitability among AWS and Microsoft:** we have found that AWS has earned persistently high profits, earning around 40% return on capital employed (ROCE) each year since 2018. This is significantly higher than our estimate of the weighted average cost of capital (WACC) of 9-13% for cloud activities. We estimate that Azure’s ROCE increased since 2018 and is now above our estimate of WACC. In comparison, despite strong revenue growth, Google Cloud has yet to make a profit. Similarly, EBIT margins for other smaller providers we looked at, Alibaba Cloud and DigitalOcean, were largely negative in recent years.

6.14 Whilst high profitability, a concentrated market structure and limited levels of switching are not necessarily concerns in themselves, they can indicate that providers are not facing sufficiently strong competitive constraints. Our provisional view is that in combination with

⁷⁰⁰ Context Consulting research report, slide 131.

⁷⁰¹ Context Consulting research data tables, Q47.

⁷⁰² 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 7.

the barriers to switching and multi-cloud we have identified in Section 5, these outcomes are consistent with a lack of effective competition.

Our provisional assessment of barriers to switching and multi-cloud contribute to our view that the market is not working well

- 6.15 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.
- 6.16 First, we have found that there are inherent technical barriers to switching cloud infrastructure services. As detailed in Section 5, Lack of technical interoperability and portability, there will always be some technical effort required of customers to switch between clouds. This is common to many IT markets – it stems from technical differences in the way different providers have designed their systems and may be beneficial for customers, e.g. by allowing providers to introduce certain innovations. These differences result in 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.
- 6.17 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 environment.
- 6.18 Second, we have found that there are certain practices by the hyperscalers that may further raise barriers to switching and multi-sourcing:
- a) **Egress fees:** The hyperscalers charge customers egress fees to transfer data out of their cloud infrastructure and our analysis indicates that the hyperscalers have set the level of egress fees significantly above the cost of transferring data. Egress fees are likely to be a significant barrier to customers using integrated multi-cloud, particularly where large volumes of data need to be transferred between clouds. This is likely to affect many customers given the wide uptake of data analytics and data management services. Egress fees can also be a material barrier to switching for some, particularly where it creates a significant one-off cost for customers to move their data from one provider to another. This cost may be particularly high for customers that process large amounts of data and have been ingesting data into cloud for many years. We also believe this cost is likely to increase for many customers in the future, as the trend is for customers to create and use more data within cloud over time.

- b) **Restricting interoperability with third-party services:** We have seen examples of AWS and Microsoft putting in place restrictions that appear to limit the interoperability of some of their services with the services of other cloud providers by not openly sharing information such as APIs and protocols. Where this is the case, cloud services can only be used within their respective clouds (e.g. Cloud Provider A database service can be used with Cloud Provider A storage service, but not with Cloud Provider B storage service). This may prevent customers from switching and building integrated multi-cloud architectures where they want to include such services within their workloads. In addition, we have seen indications that AWS and Microsoft combine proprietary technologies with open-source software and open standards for some of their first-party services. This may increase the reconfiguration efforts required to switch and multi-cloud. Some ISVs have also raised concerns that the hyperscalers favour their own first-party PaaS products where they compete with third-party ISV products. These concerns mainly relate to the level of interoperability the hyperscalers provide to ISVs and how this makes it less attractive for customers to take-up or switch to ISV products.
- c) **Committed spend discounts:** All three hyperscalers use committed spend discounts, 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. 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.

6.19 What we have heard from customers indicates that each of these practices individually may have a material influence on their behaviour by making it less attractive to switch and multi-source. For example, customers told us that egress fees are the biggest single challenge to using multi-cloud (45% of respondents)⁷⁰³ and there is evidence of customers seeking alternative cloud architectures to minimise or avoid egress fees. We have also heard that committed spend discounts are an important commercial consideration for customers. In particular, they say these discounting incentives encourage them to purchase most of their cloud services from the same cloud provider.

6.20 However, there is the potential for these practices to have an even greater impact when they combine with each other. 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 their spend between two providers causes them to miss out on their volume commitments. These incentives can ultimately result in 'path dependence',

⁷⁰³ Context Consulting research data tables, Q31; and Context Consulting research report, slide 79.

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.

- 6.21 Furthermore, by incentivising customers to build all or most of their cloud architecture in one ecosystem, these practices can add to the inherent technical barriers and make switching existing workloads more difficult. This is because switching cloud solutions that are built on several products within an ecosystem will likely involve porting all of these services from one provider to another. This is likely to lead to higher costs for customers to re-engineer all of their applications and data, compared to a scenario where they are switching only some of their services (i.e. under a multi-cloud approach).⁷⁰⁴ The need to move all services may further discourage customers to switch in cases where they would lose some 'must-have' services that are not currently available on the target cloud.
- 6.22 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 based on cloud-neutral design principles, e.g. by using container services, using 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 (e.g. smaller start-ups).⁷⁰⁵ Our evidence suggests that the vast majority of customers are likely to have more complex needs, e.g. our market research 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.⁷⁰⁶
- 6.23 Our provisional view is that these barriers are likely to result in a material share of customers facing limited possibilities to switch some (i.e. multi-cloud) or all of their workloads once they have chosen their primary provider. This is consistent with our market research which 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%), with concerns about egress fees and interoperability also ranking highly (55% and 52% respectively).⁷⁰⁷ Our evidence indicates that more mature and larger organisations may be particularly affected by technical barriers to switching.⁷⁰⁸ This may be because these customers are more likely to have large numbers of applications and/or use

⁷⁰⁴ In our market research the most cited barrier to completely switching the IaaS/PaaS provider was time and cost of making the change (43%). See: Context Consulting research data tables, Q52.

⁷⁰⁵ 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.

⁷⁰⁶ 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 raises 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..

⁷⁰⁷ Context Consulting research data tables, Q63; and Context Consulting research report, slide 131.

⁷⁰⁸ 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.

various proprietary services offered by their cloud providers, which add to the complexity of switching cloud provider.

- 6.24 In the next section, we discuss the impact of these high barriers to switching and multi-cloud on customer welfare and the potential risks to competition in future.

Potential harms to cloud infrastructure customers

Potential harms to customers that face significant barriers to switching

- 6.25 Based on the evidence we have received, we have concerns that a material number of customers face significant barriers to switching and are already experiencing harm as a result. We are most concerned in relation to AWS and Microsoft Azure, given they are the market leaders, earn profits above the estimated cost of capital and undertake all the barriers to switching and multi-cloud in some form.
- 6.26 Our provisional view is that many existing customers have a limited ability to credibly threaten to switch all or some of the workloads when they renew their contracts with their existing cloud provider. In this case, the cloud provider can exercise a degree of market power over these customers, which can lead to harm if providers are able to raise prices upon contract renewals (e.g. by reducing the discounts offered to these customers).⁷⁰⁹ They may also be harmed if there is a better product on offer by a competitor (for example, at a lower price point, or a more innovative solution) and they cannot switch their existing workloads.
- 6.27 In principle, there is scope for rivals to compete over incremental workloads that customers may wish to deploy on the cloud. However, where the barriers to using a rival provider for existing workloads are sufficiently high, a customer may have limited bargaining power over these incremental workloads:
- a) If an incremental workload needs to integrate with existing solutions on a customer's primary cloud provider, then it is likely to be significantly more costly to use a rival provider. This can lower the threat of this customer deploying the incremental workloads on a rival's cloud infrastructure, lessening the competitive constraint on their primary provider and leading to higher prices for such incremental workloads.
 - b) Where customers have financial incentives to procure incremental workloads from their current cloud provider (e.g. to avoid paying egress fees or due to the threat of losing their discount over existing cloud usage), this can lead them to settle for lower quality alternatives from their existing provider rather than best in breed rival products.
- 6.28 In summary, our provisional view is that cloud providers are likely to exert some market power in respect of the existing and incremental workloads of a material number of

⁷⁰⁹ For example, upon contract renewals, providers could raise prices to established customers by lowering the discount offered to these customers, while continuing to provide new customers with heavily discounted offerings to attract them onto their cloud.

existing customers. This could lead to negative outcomes for these customers along a number of parameters, such as pricing, quality of service and choice of offering.

- 6.29 While this is a growing market and there will always be new customers, the wave of migration from on-premise IT solutions to the cloud will diminish. Over time, as more and more customers transition from being new customers to existing customers, the number of customers who face material barriers to switch and multi-cloud will increase and with that, the scope for harm increases.

Customers are unlikely to fully offset the harms from lock-in when making the initial purchase

- 6.30 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. In particular, customers may use a variety of approaches when making their original purchases as a means of mitigating against future harms. However, we consider that this is unlikely to protect customers sufficiently once they have made their initial choice of cloud provider.
- 6.31 The barriers to switching we identified may not be front of mind when some customers first move to the cloud, because the focus may be on shorter-term objectives such as maximising savings from moving legacy software into cloud and doing so quickly. As such, they may be less inclined to make more expensive choices today that allow for greater flexibility in the future when the benefits of doing so are not always apparent when first moving to the cloud. Further, only about a third of customers have both a formal and comprehensive cloud strategy, which suggests that many customers may not be planning for their future cloud needs.⁷¹⁰
- 6.32 We have also heard that in some cases, AWS and Microsoft make tweaks to their services built on open standards and that customers may not always be aware that these are different to the original open-source version. As a result, customers may not realise when they first move into the cloud that they will still need to incur substantial costs in future to switch all or parts of the solutions they have built on these services.
- 6.33 With respect to contracting solutions, customers could notionally use their initial bargaining power to negotiate long-term contracts that safeguard their position going forward. Alternatively, they could use their bargaining power to extract a good deal in their initial contract that would offset higher prices in future contracts. However, we have seen evidence that there are limitations to the extent to which customers can negotiate terms with their providers (see Section 4). In addition, some customers may have limited bargaining power when they first move into the cloud, for example, if they need services that only certain cloud providers offer. For example, for broadcasting and video processing customers, there would have been few alternatives to AWS's Elemental at the time of first migrating to the cloud. This could limit their ability to negotiate better prices upfront.

⁷¹⁰ Context Consulting research data tables, Q20.

- 6.34 Even where customers do have bargaining power initially, there are factors that can make it difficult to use this position to extract favourable terms that would offset the harms of having limited options to switch when renewing their contract in the future.
- 6.35 There are several reasons why customer demand is not fully predictable when they first move into cloud, which makes it difficult for customers to protect themselves fully through contracting solutions when they first move into the cloud:
- a) The market research showed that in one in three cases, users found it difficult to accurately predict the future costs of their cloud computing,⁷¹¹ which suggests some find it difficult to predict their usage levels.
 - b) While most large enterprises use cloud services, a large portion of their workloads are still on-premises. It can take many years to design the migration process to move workloads onto the cloud, which means that future needs may not be known when a customer initially moves into the cloud.
 - c) Customer needs evolve as organisational needs change. Furthermore, as the market evolves and new products and capabilities are being launched, customers will inevitably look to adjust their cloud solutions to reflect the latest market developments.
- 6.36 Some customers also find it difficult to understand how providers price their services and consequently find it difficult to accurately compare offers in the market, or predict the future spend of their cloud usage. Our market research found that about half of customers are concerned about a lack of price transparency.⁷¹² For such customers, this lack of pricing transparency can weaken their bargaining power when they first move into the cloud.
- 6.37 Furthermore, the cloud services market is complex and dynamic and the purchasing of cloud solutions involves many parameters beyond price. This includes factors that we consider could vary by customer, such as ease of integration and continued engineering support. Information asymmetries inherent in contracting, together with the complexity of the services purchased, implies that it can be difficult to contract for all relevant parameters in an agreement upfront. This leaves scope for a cloud provider to extract rents from customers that face high barriers to switching through a range of metrics other than price.

Evidence from customers suggest that harms are potentially already arising today

- 6.38 The evidence we have collected as part of this market study suggests barriers to switching and multi-cloud are already creating harm for some customers today.

⁷¹¹ Context Consulting research report, slide 95.

⁷¹² Context Consulting research report, slide 131.

- a) Our market research⁷¹³ and customer engagement⁷¹⁴ found that one of the concerns for customers was de facto lock-in, and that being reliant on an ever-increasing number of interdependent services can leave them with limited ability to respond to price rises.⁷¹⁵ Our market research showed that ‘being locked-in to an individual providers’ offerings’ being a concern for about half of customers surveyed.⁷¹⁶
 - b) Our market research shows that price rises are relatively common for customers who have renegotiated a contract with their provider. It shows that 44% of respondents who renegotiated a contract reported experiencing a price rise for some or all of their services. It also indicated that such price rises can be significant, with average price rises of 20% and a median price rise of 10% among the hyperscaler customers who faced higher prices upon re-contracting.⁷¹⁷
 - c) As detailed above, we estimate that both AWS and Microsoft Azure earn profits above WACC, which indicates that they are able to maintain high prices for certain customers. AWS in particular has persistently earned high profits of around 40% ROCE each year since 2018. This compares to our estimate of WACC of 9-13%.
- 6.39 As discussed in Section 5, we have heard from [redacted] customers, in particular, that upon renegotiation of a contract, the hyperscaler may adopt a take-it-or-leave it approach to the negotiation and that some customers are being pushed to spend more to maintain their original discounts.
- 6.40 We will continue to gather evidence and views on the extent to which customers are being harmed as a result of facing significant barriers to switching.

There is a risk that the extent of competition may deteriorate in the future

- 6.41 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.
- 6.42 In contrast, barriers to entry and expansion were more material for those that entered later. Rival providers such as Microsoft and Google have had to invest significantly ahead of customer acquisition. Unlike AWS, they had to challenge an incumbent who had an established product range and data centre footprint and benefited most from economies of scale and network effects. Rival cloud providers also faced greater barriers to acquire

⁷¹³ Context Consulting research report slides 56, 59 and 110.

⁷¹⁴ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁷¹⁵ [redacted] CFI response page 5.

⁷¹⁶ Context Consulting research report, slide 131.

⁷¹⁷ Context Consulting research report, slide 101.

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.

- 6.43 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 who can 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 their 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, it appears to have benefited most from its leadership position in adjacent software markets.

There is a risk that the market becomes concentrated with a few ecosystem providers only

- 6.44 Looking ahead, we consider that there is the potential for the market to become more concentrated around the market leaders, with scope for the harm we observe today to increase and 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:
- a) 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.
 - b) We anticipate that there will be fewer new customers to vie for once most existing businesses have moved their workloads into the cloud. Given the factors above, many new customers are likely to favour the ecosystems of AWS or Microsoft, and with dwindling numbers of new customers it could become even harder for smaller cloud providers to grow their customer base and close the gap to the market leaders.
 - c) The ability of smaller cloud providers to gain scale by competing for the business of AWS or Microsoft's customers (either their existing or incremental workloads) is also likely to be limited. 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 unnecessary limitations on interoperability, artificially increase the hurdles for smaller providers to compete for customers' workloads. We have also seen evidence that some customers switch away from smaller cloud providers to avoid egress fees. This suggests that smaller cloud providers are already limited in their ability to compete for existing customers' workloads, and the number of customers affected by this is likely to increase in future

as customers further expand their cloud infrastructure with a primary provider and more customers become locked-in.

- d) Market shares have remained concentrated towards the market leaders in recent years. While Google is often positioned as the closest challenger and has been gaining share, it remains far behind in terms of size and is yet to make a profit. The other smaller cloud providers remain further behind.
- 6.45 If these factors combine to make it 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. This effect may be aggravated if difficulties to grow their customer base also makes it harder for smaller cloud providers to attract ISVs into their ecosystems.
- 6.46 This could have long lasting impacts if it leads the market to become more concentrated towards a few ecosystems, with barriers to switching and multi-cloud allowing the market leaders to entrench their positions and avoid competing vigorously with each other. 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 rivals like Google. [X] Microsoft's Azure returns [X] 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 rivals like Google.
- 6.47 Ultimately whether we arrive at a market concentrated around the two market leaders only or a concentrated market with limited competition between a few ecosystem providers, we consider that this has the potential to lead to worse outcomes compared to what is observed today, as we discuss below.

Market concentration towards a few ecosystems may further inhibit effective competition

- 6.48 Should we arrive at a more concentrated market where smaller cloud providers pose a weaker constraint on the ecosystems of the market leaders, this has the potential to dampen competition and lead to worse outcomes for customers going forward. This is both in terms of innovation and pricing outcomes.
- 6.49 Competition can drive innovation. This is particularly important in a dynamic market like cloud infrastructure services, which is characterised by significant scope for continuous improvement of existing products and the development of services that cater to new use cases. We observe that ISVs and smaller cloud providers play an important role in developing innovative solutions to attract new customers and new workloads, with the market leaders developing their own solutions in response.
- 6.50 We consider that a more concentrated market may chill innovation if smaller cloud providers are less able to compete for new customers and new workloads, therefore reducing the incentives of market leaders to invest and improve their product offering in response. Furthermore, as the market matures and there are fewer new customers to

compete for, the market leaders' incentive to invest in innovation may be weaker if there are fewer new customers to monetise that investment.

- 6.51 Weaker competitive constraints on ecosystem providers from the smaller cloud providers may also dampen price competition. To the extent that these ecosystem providers have limited incentives to chase each other's existing customer base, this may leave customers in a weaker bargaining position and ultimately lead to higher prices.

A market centred on a few ecosystems could distort competition at the PaaS layer

- 6.52 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 self-preference their own services where those compete directly with ISVs.
- 6.53 This type of self-preferencing 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.
- 6.54 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.

Question 6.1 Do you agree with our assessment of how well competition is working in cloud infrastructure and what are the potential implications of a lack of competition?

7. Microsoft's licensing practices

- 7.1 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. The concerns stem from an adjacent market outside the scope of our study. 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.⁷¹⁸
- 7.2 The concerns are set out in a number of materials and by a number of parties including the trade association CISPE,⁷¹⁹ (and the report by Professor Frédéric Jenny which they commissioned (the Jenny report)),⁷²⁰ the cloud providers [redacted]⁷²¹ and [redacted].⁷²² In 2022, CISPE filed a formal complaint against Microsoft with the European Commission relating to these practices.⁷²³ In March 2022, OVHcloud, Aruba, and the Danish Cloud Community filed a complaint with the European Commission against Microsoft. At the time of writing, there are press reports which indicate that these cloud providers are in the process of resolving their concerns with Microsoft. A settlement could result in the withdrawal of their complaint to the European Commission.⁷²⁴
- 7.3 The purpose of this section is to summarise what we have heard and set out at briefly in principle what the implications might be for cloud infrastructure competition. We are not undertaking an assessment of the conduct, nor do we make any findings in relation to the concerns themselves.
- 7.4 The section below focusses on Microsoft. There are also submissions which raise similar issues related to Oracle.⁷²⁵ 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

⁷¹⁸ Microsoft response to the s.174 notice dated 21 October 2022, question 33g, [redacted]; and Microsoft response to Ofcom's follow-up questions dated 27 January 2023 in relation to Microsoft's response to Ofcom's s.174 request of 21 October 2022.

⁷¹⁹ CISPE website, 9 November 2022. [Executive Summary of CISPE Complaint against Microsoft](#) [accessed 10 February 2023].

⁷²⁰ Jenny, Frédéric, October 2021. [Cloud Infrastructure Services: An analysis of potentially anti-competitive practices](#) [accessed 10 February 2023].

⁷²¹ [redacted] response to the CFI, [redacted].

⁷²² [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted].

⁷²³ CISPE, 9 November 2022, Complaint concerning Microsoft's abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version.

⁷²⁴ See: Microsoft in talks with OVH, others to settle EU cloud services antitrust complaint, 28 March 2023, [MLex](#) [accessed 29/03/2023]; and Microsoft offers to change cloud practices to ward off EU antitrust probe, 28 March 2023, [Reuters](#) [accessed 29/03/2023].

⁷²⁵ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], page [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], page [redacted]. See also: Jenny, Frédéric, October 2021. [Cloud Infrastructure Services: An analysis of potentially anti-competitive practices](#) [accessed 10 February 2023].

conduct has the potential to have a greater effect on competition in the cloud infrastructure market.

Summary of submissions received

- 7.5 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 expensive⁷²⁶ for customers when using Microsoft’s licensed products on third-party clouds, or other disadvantages⁷²⁷ (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.⁷²⁸
- 7.6 The submissions focus on a set of enterprise 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 enterprise software products.
- 7.7 We note that Microsoft’s software licensing practices are complex. Our summary of the practices reflects Ofcom’s provisional 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

- 7.8 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-premises software licences give them such rights, or whether they may need to repurchase their licenses or pay for additional usage rights.
- 7.9 We understand from submissions that Microsoft previously did not charge customers for deploying their software on cloud infrastructure, regardless of the cloud provider.

⁷²⁶ See, for example: CISPE, 9 November 2022, Complaint concerning Microsoft’s abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 32, paragraph 127; and [redacted] response to the CFI, [redacted], page [redacted].

⁷²⁷ See, for example: CISPE, 9 November 2022, Complaint concerning Microsoft’s abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 59, paragraph 270; CISPE, 9 November 2022, Complaint concerning Microsoft’s abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 40, paragraph 182; [redacted] response dated [redacted] to our follow-up questions dated [redacted] concerning the s.174 notice dated [redacted], question [redacted], paragraph [redacted]; and [redacted] response dated [redacted] to Ofcom’s follow-up questions dated [redacted] in relation to [redacted] response to Ofcom’s s.174 request of [redacted], question [redacted], page [redacted].

⁷²⁸ [redacted] response dated [redacted] to Ofcom’s follow-up questions dated [redacted] in relation to [redacted] response to Ofcom’s s.174 notice dated [redacted], question [redacted], page [redacted].

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.⁷²⁹

- 7.10 These changes require customers with on-premises software licences to pay additional fees to repurchase their licences for some Microsoft products⁷³⁰ if they want to use them on certain cloud providers' infrastructure, despite already owning a licence for the use of the software on premises.⁷³¹ We have heard that this amounts to paying for the same service twice.⁷³² We have received submissions that, in some cases, Azure customers do not have to pay these additional fees.⁷³³
- 7.11 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.⁷³⁴ 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.⁷³⁵ 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.⁷³⁶ Products such as Windows Server, SQL Server, and the Microsoft Office Desktop Applications are available through the SPLA.⁷³⁷ One cloud provider ([redacted]) 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.⁷³⁸
- 7.12 We have heard from competing cloud providers that purchasing some Microsoft products via the SPLA can be more expensive compared to Microsoft products hosted on Azure.⁷³⁹ We have received examples to show this. 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 ([redacted]) compared to the price that Microsoft was able to

⁷²⁹ [redacted] response to the CFI, [redacted], page [redacted]; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], paragraph [redacted].

⁷³⁰ Some of the products that customers can be required to repurchase licences for include the Microsoft Office productivity apps, Windows Desktop, and Windows Server. See: [redacted] response to the CFI, [redacted], page [redacted].

⁷³¹ [redacted] response to the CFI, [redacted], page [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], page [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], paragraph [redacted]; and CISPE, 9 November 2022, Complaint concerning Microsoft's abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 40, paragraph 180.

⁷³² [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], page [redacted]; and [redacted] response to the CFI, [redacted], page [redacted].

⁷³³ [redacted] response to the CFI, [redacted], page [redacted].

⁷³⁴ [redacted] response to the CFI, [redacted], page [redacted].

⁷³⁵ [redacted] response to the CFI, [redacted], page [redacted].

⁷³⁶ [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], question [redacted], page [redacted].

⁷³⁷ See: Microsoft document, October 2021. Services Provider Use Rights. Available for download at: [Licensing Documents \(microsoft.com\)](https://www.microsoft.com/licensing/documents) [accessed 14/03/2023].

⁷³⁸ [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], question [redacted], page [redacted].

⁷³⁹ See: [redacted] response to the CFI, [redacted], page [redacted]; [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], page [redacted]; and [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], question [redacted], page [redacted].

offer for using Azure instead.⁷⁴⁰ In this example, the competing cloud provider ([redacted]) told us that it could not match the price Microsoft was able to offer, despite discounting other infrastructure costs significantly.⁷⁴¹

- 7.13 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. In an example submitted to us, if a potential customer had chosen to move their legacy on-premise Windows Server and SQL Server footprint⁷⁴² to a competing cloud provider's cloud infrastructure ([redacted]) then over 70% of their total contract costs per year would have been for extra licensing fees arising from Microsoft's software licensing practices.⁷⁴³
- 7.14 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.⁷⁴⁴ We have heard that in some cases SPLA prices increased by 10-15% in 2018 and 2019,⁷⁴⁵ and for some providers their SPLA prices have increased by over 50%.⁷⁴⁶ One competing cloud provider ([redacted]) told us that they 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.⁷⁴⁷ However, another competing cloud provider ([redacted]) submitted that they have historically absorbed these cost increases and have not raised prices for customers after Microsoft raised their licensing fees under the SPLA.⁷⁴⁸
- 7.15 We received a submission that as a result of Microsoft's licensing practices, only a small proportion of a competing cloud provider's ([redacted]) customers run Microsoft products on their cloud infrastructure: they estimate that only <0.5% of their compute service usage

⁷⁴⁰ In this example, the customer was looking to move their Windows Server-based on-premises infrastructure to the cloud. The example explains that the cost differential was a Microsoft's licensing policy, which restricted this potential customer from using their existing Windows Server licenses on the competing cloud provider's cloud infrastructure ([redacted]). See: [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], [redacted], page [redacted].

⁷⁴¹ [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], [redacted], page [redacted].

⁷⁴² [redacted] presentation to Ofcom, [redacted], slide [redacted].

⁷⁴³ [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], question [redacted], page [redacted].

⁷⁴⁴ [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], page [redacted]; and CISPE, 9 November 2022, Complaint concerning Microsoft's abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 36, paragraph 157.

⁷⁴⁵ See: Jenny, Frédéric. October 2021. [Cloud Infrastructure Services: An analysis of potentially anti-competitive practices](#) [accessed 10 February 2023], page 37, paragraph 71; and CISPE, 9 November 2022, Complaint concerning Microsoft's abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 36, paragraph 157.

⁷⁴⁶ [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], question [redacted], page [redacted]; and [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], question [redacted], page [redacted], paragraph [redacted].

⁷⁴⁷ [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], question [redacted], page [redacted].

⁷⁴⁸ [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], question [redacted], page [redacted], paragraph [redacted].

runs on Windows 10/11, and an estimated <5% of their compute service usage runs on Windows Server.⁷⁴⁹

- 7.16 Our engagement with customers also suggests that Microsoft’s licensing policies are a factor that impacts on customer choice as discussed in Section 5, importance of existing relationships in adjacent software markets.
- 7.17 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.⁷⁵⁰ 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. Customers are eligible to do so if they purchase a ‘Software Assurance’ subscription⁷⁵¹ with their existing on-premises Microsoft software licences.⁷⁵²
- 7.18 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.⁷⁵³ We have heard that this means customers are not able to deploy their existing on-premises licences for certain software on Listed Provider cloud infrastructure,⁷⁵⁴ and in these cases are required to pay additional fees to licence these products again.⁷⁵⁵
- 7.19 CISPE has submitted that these changes have failed to address industry concerns.⁷⁵⁶ Listed Providers remain excluded from these changes⁷⁵⁷ and CISPE explains it is not clear who exactly will qualify to avoid the additional licensing fees that can be required of Listed Providers.⁷⁵⁸ We have also received concerns from stakeholders that there is no restriction

⁷⁴⁹ [redacted] response dated [redacted] to Ofcom’s follow-up questions dated [redacted] in relation to [redacted] response to Ofcom’s s.174 notice dated [redacted], question [redacted], page [redacted].

⁷⁵⁰ Microsoft website, 18 May 2022. [Microsoft responds to European Cloud Provider feedback with new programs and principles - EU Policy Blog](#) [accessed 20 March 2023].

⁷⁵¹ 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 14/03/2023].

⁷⁵² CISPE, 9 November 2022, Complaint concerning Microsoft’s abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 40, paragraph 180-183; and [redacted] response dated [redacted] to Ofcom’s follow-up questions dated [redacted] in relation to [redacted] response to Ofcom’s s.174 notice dated [redacted], question [redacted], page [redacted].

⁷⁵³ CISPE, 9 November 2022, Complaint concerning Microsoft’s abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 40, paragraph 180-183; [redacted] response dated [redacted] to Ofcom’s follow-up questions dated [redacted] in relation to [redacted] response to Ofcom’s s.174 notice dated [redacted], question [redacted], page [redacted]. See: Microsoft document. <https://aka.ms/FlexibleVirtualizationBenefitGuide> [accessed 28/03/2023].

⁷⁵⁴ [redacted] response dated [redacted] to our follow-up questions dated [redacted] concerning the s.174 notice dated [redacted], question [redacted], paragraph [redacted]; and [redacted] response dated [redacted] to Ofcom’s follow-up questions dated [redacted] in relation to [redacted] response to Ofcom’s s.174 notice dated [redacted], question [redacted], page [redacted].

⁷⁵⁵ [redacted] response to the CFI, [redacted], page [redacted].

⁷⁵⁶ CISPE, 9 November 2022, Complaint concerning Microsoft’s abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 56-8.

⁷⁵⁷ CISPE, 9 November 2022, Complaint concerning Microsoft’s abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 40, paragraph 183.

⁷⁵⁸ CISPE, 9 November 2022, Complaint concerning Microsoft’s abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 40, paragraph 182.

on Microsoft adding other competing cloud providers to the list of Listed Providers in the future.⁷⁵⁹

- 7.20 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. 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.⁷⁶⁰
- 7.21 One 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”.⁷⁶¹ Microsoft offers the Azure Hybrid Benefit to customers moving their Windows or SQL Server workloads to Azure.⁷⁶² Microsoft Azure blog 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 infrastructure. You will need to repurchase your Windows Server license on other providers’ clouds”.⁷⁶³
- 7.22 We have also heard that for Windows 10/11, which is not available through SPLA, customers using certain competing cloud providers’ cloud infrastructure are required to pay additional fees, and at times these fees may not apply to Azure customers in the same way.⁷⁶⁴ The submissions explain that purchasing a user licence for Virtual Desktop Access (VDA) – used by customers to license Windows for use with virtual machines – is an additional cost for certain customers that in some cases does not apply when the customer is running Windows 10/11 on Azure. A competing cloud provider ([redacted]) submitted that this restriction can result in a significantly higher annual per user cost for its customers.⁷⁶⁵

Alleged practices – non-cost issues for users of rival cloud infrastructure

- 7.23 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.
- 7.24 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 customer.⁷⁶⁶

⁷⁵⁹ CISPE, 9 November 2022, Complaint concerning Microsoft’s abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 58, paragraph 263; and [redacted] response dated [redacted] to the s.174 notice dated [redacted], question [redacted], page [redacted], paragraph [redacted].

⁷⁶⁰ [redacted] response to the CFI, [redacted], page [redacted].

⁷⁶¹ [redacted] response to the CFI, [redacted], page [redacted].

⁷⁶² See: Microsoft Azure Website. [Azure Hybrid Benefit - Hybrid Cost Calculator | Microsoft Azure](#) [accessed 23/03/2023].

⁷⁶³ See: Microsoft Azure blog, 26 January 2022. [Save big by using your on-premises licenses on Azure | Azure Blog and Updates | Microsoft Azure](#) [accessed 02/02/2023].

⁷⁶⁴ [redacted] response dated [redacted] to Ofcom’s follow-up questions dated [redacted] in relation to [redacted] response to Ofcom’s s.174 notice dated [redacted], question [redacted], page [redacted]; and [redacted] response dated [redacted] to our follow-up questions dated [redacted] concerning the s.174 notice dated [redacted], question [redacted], paragraph [redacted].

⁷⁶⁵ [redacted] response dated [redacted] to our follow-up questions dated [redacted] concerning the s.174 notice dated [redacted], question [redacted], paragraph [redacted].

⁷⁶⁶ [redacted] response to the CFI, [redacted], page [redacted].

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,⁷⁶⁷ Windows Server⁷⁶⁸ and Windows Desktop⁷⁶⁹ 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.⁷⁷⁰
- b) We have also heard that customers cannot use the Microsoft 365 Apps on certain non-Microsoft cloud infrastructure.⁷⁷¹

Microsoft's response to summary regarding its software licensing practices

7.25 Microsoft disputes the veracity of the submissions.⁷⁷² When given summary details of the submissions and the opportunity to comment, Microsoft has stated that although Ofcom have 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.⁷⁷³

Relevance for competition in cloud infrastructure

7.26 The nature of the submissions that we have received is that Microsoft engages in a number of behaviours in relation to its enterprise software products which encourage customers to use this software on Azure rather than on competitors' infrastructure. The submissions argue that this conduct leads to an adverse impact on competition in cloud infrastructure.

7.27 In principle, Microsoft's behaviour in relation to its enterprise 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

⁷⁶⁷ 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: [redacted] response dated [redacted] to our follow-up questions dated [redacted] concerning the s.174 notice dated [redacted], question [redacted], paragraph [redacted].

⁷⁶⁸ 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: [redacted] response the CFI, [redacted], page [redacted].

⁷⁶⁹ CISPE have also explained that only Azure customers can run the more CPU-efficient Windows 11 multi-session. See: CISPE, 9 November 2022, Complaint concerning Microsoft's abusive practices on the market for cloud services in violation of Article 102 TFEU and proposals to remedy identified abuses, Non-confidential version, page 60, paragraph 270.

⁷⁷⁰ [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], [redacted], page [redacted].

⁷⁷¹ [redacted] response dated [redacted] to our follow-up questions dated [redacted] concerning the s.174 notice dated [redacted], question [redacted], paragraph [redacted]; and [redacted] response dated [redacted] to Ofcom's follow-up questions dated [redacted] in relation to [redacted] response to Ofcom's s.174 notice dated [redacted], question [redacted], page [redacted].

⁷⁷² Microsoft response to the s.174 notice dated 21 October 2022, question 33g, [redacted]; and Microsoft response to Ofcom's follow-up questions dated 27 January 2023 in relation to Microsoft's response to Ofcom's s.174 notice dated 21 October 2022.

⁷⁷³ [redacted].

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 6.

- 7.28 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 alternatives) 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.

Next steps

- 7.29 In our market study we have not carried out an assessment of the allegations regarding Microsoft's licensing practices. This is because (i) they, in part, concern Microsoft's position in enterprise software, any detailed consideration of which is outside the scope of our study, (ii) given the importance and complexity of the issues, we do not think it is appropriate to do this in the study, and (iii) we understand that the issues are under active consideration in other jurisdictions.
- 7.30 We do not anticipate considering these issues further in the remainder of our study. Ofcom and the CMA will consider the most appropriate way forward on these issues.

8. Overview of potential interventions

- 8.1 In Sections 5 and 6 of this interim report, we have set out our provisional conclusions on competition in cloud infrastructure services. We identified a number of reasons why the market for cloud infrastructure services may not be working well for customers and ultimately UK consumers. These are both inherent features of the market, and specific practices undertaken by the hyperscalers, in particular AWS and Microsoft.
- 8.2 This section provides an overview of the types of interventions that in our provisional view may address the potential competition concerns identified in Section 6.
- 8.3 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. We highlight our initial views on these points in relation to the potential interventions set out below, and invite feedback to support our further assessment in the second half of the study. We would also welcome feedback on any interventions that we may not have considered.
- 8.4 We have not, at this stage, sought to determine whether any individual intervention would be justified, and therefore we are not making recommendations or advocating for any specific interventions at this time. Instead, this section sets out a high-level overview of the potential merits, risks and challenges associated with the interventions we have identified as provisionally credible and effective. 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.
- 8.5 We are also aware of related regulatory developments, both in the UK with the proposed Digital Markets Regime and internationally, particularly in Europe, that have the potential to impact the competitive landscape for cloud infrastructure services. When considering potential interventions in the UK, we are cognisant of proposals elsewhere, especially given the global nature of cloud infrastructure services. Again, at this stage in our interim report, a definitive assessment is not necessary but we would invite stakeholder views on how this interaction may play out.

Improving competitive outcomes by reducing barriers to multi-cloud and switching

- 8.6 It is our provisional view that whilst there are positive signs of competition in cloud infrastructure services at present, there are a number of barriers to effective competition. We see indicators of harm today and risks for this to become worse in future, with barriers to switching and multi-cloud at the core of our concerns.
- 8.7 Our interim position is that reducing barriers to switching and multi-cloud (particularly integrated multi-cloud) would improve customer outcomes now and reduce the risk of competition deteriorating in the future. In the short-term, it would strengthen the threat of existing customers' moving current and future workloads to rival clouds, which in turn

would strengthen their bargaining power as they renegotiate their contracts. This should result in lower prices and the ability for customers to build solutions combining best-in-breed cloud products that best satisfy their needs.

- 8.8 Our provisional 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 Section 6, 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. We are concerned that this could weaken the competitive constraints smaller cloud providers pose on the market leaders, with barriers to switching and multi-cloud allowing the market leaders 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.
- 8.9 We think that the above 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 artificial 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.
- 8.10 At this interim stage, we have not identified the need for specific interventions relating 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

- 8.11 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 Section 5. We invite feedback on the effectiveness and potential risks of these interventions.

Data egress fees

- 8.12 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. However, some cloud providers charge customers for transferring that data out of their cloud infrastructure (egress fees). Many customers have told us that egress fees have an impact

on their choices, incentivising the purchase of multiple cloud services through the same provider and limiting their ability to operate a multi-cloud architecture, particularly an integrated one.

- 8.13 We have provisionally identified three key approaches that could be considered to address the competition concern in the UK market raised by providers charging egress fees.

Equalise egress fees with other charges

- 8.14 One approach to mitigating the effect of egress fees on incentives to switch and multi-cloud 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, fostering competition based on quality of services and the price of using those services, rather than the price of transferring data to them.
- 8.15 Determining exactly which other data transfer charges to equalise egress fees against would be complex. Charges vary based on factors such as the location of originating and destination data centres and type of infrastructure used to transfer the data (e.g. private network or public internet). 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”

- 8.16 A price control that restricts egress fees to “at cost” charges is likely to reduce the price of egress from current levels, making 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. This would mean customers who are transferring the most data are still charged in line with their usage, rather than being subsidised by others.
- 8.17 With a price control at cost, while there is a risk of a potential waterbed effect on other prices, the evidence we have at this stage suggests that this risk is likely to be low and the scope of any impact may be limited. This is because [§<] and there are indications that the hyperscalers are setting the level of egress fees above the cost of transferring data.
- 8.18 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. In relation to the underlying costs of egress data transfer, while our provisional assessment is that egress fees are not directly linked to costs (see the discussion on egress in Section 5), we have only seen limited information on this issue. We will undertake further work to understand this better during the second half of the study.

Prevent providers from charging for data egress

- 8.19 The most straightforward way of designing an intervention on egress fees is to prevent cloud providers from charging them. 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. Measures to prevent egress fees are currently being considered in other jurisdictions, for example by the Dutch competition authority⁷⁷⁴ and at EU level within the draft Data Act.⁷⁷⁵ Removing egress fees could have the largest impact on enabling the feasibility of multi-cloud deployments and switching.
- 8.20 Some smaller providers, such as OVHCloud and Oracle, set significantly lower egress fees. This suggests that use of alternative pricing structures is unlikely to affect investment incentives, as cloud providers will continue to compete for customers based on the price and quality of their products and services. In fact, by increasing the competitive constraint from rival providers, the incentive to invest may increase.
- 8.21 There may be adverse impacts from preventing any egress fees being charged. Providers would need to recover the costs of providing egress from other charges, which could lead to a waterbed effect on these prices. Beyond that, removing egress fees could also affect provider's incentives in relation to other prices, such as internal data transfer fees, which may in turn increase the size of any potential waterbed effect.
- 8.22 The intention of the remedy would be to make data egress more feasible. This would increase the ability of customers to multi-cloud and switch, increasing competition between cloud providers. However, 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.

Question 8.1: Do you agree that egress fees are an area of potential intervention? How might such an intervention be approached?

Technical interoperability and portability

- 8.23 In response to our CFI, we received several recommendations on different measures related to technical interoperability and portability of cloud infrastructure services.⁷⁷⁶ At this stage, we have undertaken only a preliminary assessment of the different interventions which may be considered, focusing on their benefits, feasibility and any risks

⁷⁷⁴ ACM, 2022. [Market Study Cloud services](#) [accessed 23 March 2023].

⁷⁷⁵ European Commission, 2022. [Data Act: Commission proposes measures for a fair and innovative data economy](#) [accessed 23 March 2023].

⁷⁷⁶ Vodafone CFI response; Cloudflare CFI response; [3<].

of unintended consequences. We would welcome stakeholder views on these, which we have broadly grouped into three categories:

- a) Requirements for providers to be more transparent about the interoperability of their services
- b) Requirements for the hyperscalers to make their own services easier to interoperate with third-party services
- c) Requirements which aim to change how cloud infrastructure services are designed to increase the degree of standardisation

Requirements for providers to be more transparent about the interoperability of their services

- 8.24 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. Some stakeholders have also alleged that AWS and Microsoft make inaccurate claims about their cloud infrastructure services being highly compatible with competing open-source software. Inadequate transparency and documentation can serve as a barrier to switching and integrating multiple clouds. For example, customers may build a system which they believe to be highly interoperable, without realising that they are relying on different cloud technologies.
- 8.25 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 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 switching and multi-cloud.
- 8.26 However, we are not certain about the potential unintended consequences of 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.
- 8.27 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. This could include information on the value of cloud-agnostic services, intermediaries, adaptors,⁷⁷⁷ or other ways of increasing the ease of switching and multi-cloud.

⁷⁷⁷ An adaptor is a piece of software that intermediates the communication between two or more components that cannot directly interoperate with each other. Adaptors can act as abstraction layers that translate and bridge communication between otherwise incompatible APIs, formats and protocols.

- 8.28 However, there is already substantial information available and so the incremental impact of publishing additional advice would likely be low. It may also not be sufficient to change customer behaviour in a way that truly fosters competition: intermediaries and adaptors can increase customer costs, such as re-engineering fees, egress fees, and subscription costs for additional services.

Requirements to make first-party services easier to interoperate with third-party services

- 8.29 As set out in Section 5, our analysis to date indicates that hyperscalers, particularly AWS and Microsoft, may put in place restrictions that limit the interoperability of some of their services with the services of other cloud providers. This limits customers' ability to switch because they are not able to separate out the service they want to keep (e.g. a PaaS service) from that which they might want to switch (e.g. an IaaS service). It also limits customers' ability to integrate multiple clouds where hyperscalers' cloud services only interoperate with the cloud technology in their respective clouds.
- 8.30 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 their third-party cloud services. This would potentially distort the ability for rival ISVs to compete with the hyperscalers' first-party products.
- 8.31 A way to address these concerns would be to set requirements that effectively separate or 'unbundle' cloud providers' own first-party services into their respective elements. The aim would be to allow third-party services to interoperate with those individual elements in the same way as AWS and Microsoft currently do when providing their own 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.
- 8.32 First, a requirement for cloud providers to make the individual elements of their cloud infrastructure services available to use in conjunction with those of other providers. In practice this would mean separating out first-party IaaS and PaaS services so they can be used together with a complementary service from another provider. For example, this would allow a customer to purchase AWS PaaS service and use it in conjunction with the IaaS service of a rival provider. This would increase customer choice and require hyperscalers to facilitate interoperability with their cloud services, thereby lowering the costs of switching and integrating multiple public clouds.
- 8.33 However, the requirement to make services available might only guarantee a basic level of technical interoperability. Customers may still find integration is more difficult or quality of experience is reduced when using first-party and third-party services in combination, by comparison to using an integrated service provided solely by one of the hyperscalers.
- 8.34 This means a second requirement may be needed to ensure services from another cloud provider or ISV (for example a type of PaaS service), can interoperate with a cloud provider's services (for example storage or another type of IaaS service), in the same way

as the original cloud provider's first-party services. In practice, this could be achieved by a requirement for cloud providers to 'open up' their cloud infrastructure services (for example, by opening up the APIs and/or protocols associated with them), so they can be accessed on an equivalent basis by other cloud providers or ISVs. This would also limit the ability of cloud providers to circumvent the initial requirement by making their services available, but with only a limited degree of interoperability.

- 8.35 We also consider that such an intervention (likely alongside reductions in egress fees) could increase the demand for multi-cloud, which could incentivise market-led solutions that can facilitate interoperability between clouds. This could help address wider technical limitations to multi-cloud, which we identified in Section 5.
- 8.36 The exact nature and impact of these types of interventions would require careful consideration. For example, requirements on cloud providers to adapt or open up some technical elements of their services would lead to implementation costs. If it leads to less control over the architectural and operational decisions relating to their first-party cloud services, this intervention may also dampen cloud providers' incentives to innovate. There may also be challenges with defining the exact scope, including whether they would apply to specific services or to specific cloud providers.

Standardisation of cloud technologies

- 8.37 In Section 5 we described how the lack of standardised cloud technologies associated with interoperability, and application and data portability could be a barrier to switching and integrating multiple clouds. These technologies include APIs, protocols, workflows, programming languages, data formats or other technologies the customer applications rely on. While there are services available in the market to assist customers to mitigate the lack of standards through various means, there appears to be limitations to their effectiveness.
- 8.38 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.
- 8.39 One intervention would be to use standards to improve interoperability and portability in cloud services. 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 voluntarily approach means the incentives of hyperscalers are unlikely to align with the use of established standards unless there is significant industry pressure to do so. We have seen evidence indicating that the hyperscalers building first-party services on open standards and open APIs, but making adjustments so that the resulting first-party service is ultimately less compatible with the underlying standards.
- 8.40 Another approach would be to mandate the use of specific standards. While this intervention would guarantee broader adoption, it also comes with significant complexity and risks of unintended consequences.

- 8.41 Where standards already exist, they are often for the most basic building blocks of cloud, and therefore do not address the types of interoperability and portability challenges identified in our research to date. Designing suitable cloud standards would be challenging, and a range of issues would need to be addressed. These include deciding which services or technologies 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.
- 8.42 Mandating standards is likely to lead to significant implementation costs where they require re-engineering of cloud services and customer applications. It also carries a risk to innovation in a dynamic market such as cloud, where any standards could become outdated if they are insufficiently flexible to adapt to industry developments.

Question 8.2: Do you agree that interoperability and portability are areas of potential intervention? How might such an intervention be approached?

Committed spend discounts

- 8.43 As explained in Section 5, committed spend discounts involve a customer agreeing to spend a set amount on a cloud provider's products over a set period of time. In return, they receive a percentage discount. Such discounts are used by all three hyperscalers, especially in negotiations with the largest cloud customers – a high proportion of whom have substantial spending commitments to hyperscalers.
- 8.44 We have seen indications that hyperscalers privately negotiated committed spend discounts 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. This has the potential to dampen competition for customers' new workloads.
- 8.45 One intervention that could be considered to curb the impact of committed spend discounts on market outcomes could be prohibiting or restricting certain discount structures that may be loyalty-inducing and risk distorting competition.
- 8.46 However, discounts are generally a positive feature of markets, leading to lower prices and promoting competition between providers. There may also be some efficiency reasons for discounting structures – for example to help providers achieve scale and to motivate investment. This means that placing restrictions on discounting practices could lead to unintended consequences. For example, the bargaining power of large customers may be weakened if the structure of discounts is restricted, or if hyperscalers' incentive to invest may be dampened if they have less certainty about their future demand.

Question 8.3: Do you agree that committed spend discounts are an area of potential intervention? How might such an intervention be approached?

Challenges predicting cloud spend

- 8.47 As set out in Section 5, some customers have told us about the challenges they face when trying to understand pricing and billing processes undertaken by cloud providers, as well as issues associated with trying to predict future cloud spend. Seventy-eight per cent of respondents to our market research told us that they believed it should be easier for businesses to understand their bills and control their costs.⁷⁷⁸ We heard that the majority of firms want greater transparency around pricing, and better tools for cost monitoring and control.⁷⁷⁹
- 8.48 We are aware of industry-led approaches to addressing these concerns that could have a beneficial effect on competition. For example, [3<] acknowledged that they need to simplify their pricing structures for customers and provide sellers with appropriate material and tools to explain pricing structures.⁷⁸⁰ 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.
- 8.49 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. Nonetheless, there are likely to be practical challenges with such an approach, given the complexity and diversity of cloud services, and that it could impede innovation by providers who sell their services.

Question 8.4: Do you consider that transparency of billing could be an area of potential intervention? How might such an intervention be approached?

Skills

- 8.50 As set out in paragraphs 5.55-5.60 of 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.
- 8.51 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 switching and multi-cloud. 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.

⁷⁷⁸ Context Consulting research report, [slide 136].

⁷⁷⁹ Context Consulting research report, [slide 142].

⁷⁸⁰ [3<] response dated [3<] to s.174 dated [3<], [3<].

- 8.52 Addressing concerns around the availability of sufficiently broad skills in a technology area is always difficult and multifaceted, and potentially an area that is less well suited to regulatory intervention. We note that the Government has a number of planned actions and initiatives in relation to skills, which could generate broad benefits of relevance to cloud sectors.
- 8.53 On 6 March 2023, the Department for Science, Innovation and Technology (DSIT) published its UK Science and Technology Framework. This included ten actions that the Government plans to take in relation to critical technologies identified by the framework to support innovation in the UK.⁷⁸¹ One set of actions relates to talent and skills, which aim to support relevant UK skills as well as technical and entrepreneurial talent in STEM, digital, data, commercialisation and national security.
- 8.54 In addition, an independent report commissioned by His Majesty's Treasury and DSIT, published in March 2023, recommended that the government "support the pipeline of people with compute skills through investment in computational and digital skills, and the implementation of the Digital Strategy".⁷⁸²

Further input from stakeholders

- 8.55 We are not making a recommendation on any specific interventions at this stage. However, our provisional view is that there are credible interventions available that could address the different barriers to competition that we have identified through this market study to date. We welcome stakeholder views on these and any other potential interventions.
- 8.56 We recognise, based on our preliminary 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.
- 8.57 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 are particularly interested in feedback on how the interrelationship between measures to address concerns regarding egress fees and those regarding interoperability. For example, if egress fees were to become less of a barrier to what extent does that impact the need for regulatory actions to improve interoperability and vice versa.

⁷⁸¹ The Framework identifies the following critical technologies: artificial intelligence, engineering technology, future telecommunications, semiconductors, and quantum technologies. [The UK Science and Technology Framework](#) [accessed 16 March 2023].

⁷⁸² DSIT, 2023. [Independent Review of The Future of Compute: Final report and recommendations](#) [accessed 16 March 2023].

Question 8.5: What, if any, potential unintended consequences do you anticipate might be associated with the interventions set out above, and how might they interact with each other if implemented?

9. Next steps

- 9.1 This interim report has set out our provisional views on competition in cloud infrastructure services in the UK, following the first six months of our market study. As we work towards our final report, we want to gather more evidence to test and refine our thinking, and to design the most appropriate outcomes for our market study. Publication of this interim report for consultation is an important first step in that process.

Provisional view of outcomes

- 9.2 At the interim stage our provisional recommendation is that we refer the market for public cloud infrastructure services to the CMA to carry out a market investigation. We believe we have reasonable grounds to suspect that a feature or combination of features of a market or markets in the UK prevents, restricts, or distorts competition, in particular: egress fees, restrictions on interoperability, and committed spend. We have published a separate consultation setting out the case for a market investigation reference, which is also summarised in Section 10 of this interim report.⁷⁸³
- 9.3 We have also identified other barriers that may require intervention, in particular technical skills and transparency of billing. The impact of these issues, and the ability to design an effective intervention if required, is less certain. We welcome stakeholder views on these areas.
- 9.4 As set out in Section 7 we have heard concerns regarding Microsoft's licensing practises which we do not intend to consider further in the remainder of our market study. Ofcom and CMA will consider the appropriate way forward for these issues

Second half of the market study

- 9.5 We are consulting on the provisional findings in our interim report and the proposal to refer public cloud infrastructure services for a market investigation. We would welcome feedback from any interested parties and hope to gather views from a diverse range of stakeholders. We will assess any responses before reaching a final decision on the most appropriate outcomes for this market study.
- 9.6 A particular focus for us in the second half of the study will be to obtain further evidence on: (i) the specific issues we have identified as being the main barriers to the market working well; and (ii) the scope for such barriers to be successfully addressed through interventions.
- 9.7 In reaching our final decisions we will continue to track and further develop our understanding of the broader regulatory landscape on cloud computing, and use our interim findings to help contribute to any related ongoing work.

⁷⁸³ Ofcom, 2023. [Public cloud infrastructure services, Consultation: Proposal to make a market investigation reference.](#)

- 9.8 Internationally, of particular interest are the European Union’s Digital Markets Act and draft Data Act, which together cover questions of data portability, interoperability, and egress fees.⁷⁸⁴ In the UK, we will continue to engage with other regulators with responsibilities that relate to cloud services, such as the FCA⁷⁸⁵ and the ICO.⁷⁸⁶
- 9.9 The Digital Markets, Competition and Consumer (DMCC) Bill – expected to be introduced in the third parliamentary session – will establish a new pro-competition regime for digital markets.⁷⁸⁷ The Government intends that the DMCC Bill will empower the CMA to designate firms providing ‘digital activities’⁷⁸⁸ with Strategic Market Status (SMS) and apply binding Conduct Requirements, to manage the effects of market power, and Pro-Competitive Interventions, to tackle the root causes of their market power. Further consideration will have to be given to the question of how the new regime will apply to the firms considered in this market study. We will work with the CMA to progress that thinking as greater detail on the regime becomes available.
- 9.10 We will publish our final report by 5 October 2023. It will set out our findings and our decision on whether or not to refer public cloud infrastructure services to the CMA to carry out a market investigation.

⁷⁸⁴ See our [Call for Inputs](#), 2022.

⁷⁸⁵ For more information on the FCA’s work relating to critical third parties used by financial services firms which includes cloud services, see the FCA/Bank of England [Discussion paper on Operational resilience: Critical third parties to the UK financial sector](#) [accessed 27 March 2023].

⁷⁸⁶ 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 27 March 2023].

⁷⁸⁷ During the Autumn Statement on 18 November, Chancellor Jeremy Hunt confirmed that the Government will bring forward the Digital Markets, Competition and Consumer Bill in the third Parliamentary session. See [Autumn Statement 2022](#) [accessed 13 March 2023].

⁷⁸⁸ According to the Government’s consultation response for the pro-competition regime for digital markets, “the scope of the regime will be limited to ‘digital activities’” and “the government will work to develop a definition of the activities in scope which is clear and easy to apply.” See [A new pro-competition regime for digital markets – government response to consultation](#) [accessed 13 March 2023].

10. Our decision to consult on a market investigation reference

10.1 As per section 131B(1) EA02, we have a duty to publish our provisional view of whether we are minded to make a market investigation reference (MIR), within six months of launching our market study. Having assessed our competition concerns and the potential approaches to mitigation, we believe that we have reasonable grounds to suspect that there are features in the public cloud infrastructures services market that may have an adverse effect on competition in the UK. In particular:

- a) **Egress fees:** The hyperscalers charge customers egress fees to transfer data out of their cloud infrastructure and our analysis indicates that the hyperscalers have set the level of egress fees significantly above the cost of transferring data. Egress fees are likely to be a significant barrier to customers using integrated multi-cloud, particularly where large volumes of data need to be transferred between clouds. This is likely to affect many customers given the importance of data analytics and data management services. Egress fees can also be a material barrier to switching for some, particularly where it creates a significant one-off cost for customers to move their data from one provider to another. This cost may be particularly high for customers that process large amounts of data and have been ingesting data into cloud for many years. We also expect this cost is likely to increase for many customers in the future, as the trend is for customers to create and use more data within cloud over time.
- b) **Restrictions on interoperability and portability:** This is common to many IT markets, and it stems from technical differences in the way different providers have designed their systems. These differences result in 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. We have also seen examples of particular practices by cloud providers that appear to limit the interoperability and portability of some of their services with the services of other cloud providers by not openly sharing information such as APIs and protocols. This may prevent customers from switching and building integrated multi-cloud architectures where they want to include such services within their workloads. In addition, we have seen evidence that hyperscalers, particularly AWS and Microsoft, combine proprietary technologies with open-source software and open standards for some of their first-party services. This may increase the reconfiguration efforts required to switch and multi-cloud. Some ISVs have also raised concerns that the hyperscalers favour their own first-party PaaS products where they compete with third-party ISV products. These concerns mainly relate to the level of interoperability the hyperscalers provide to ISVs and how this makes it less attractive for customers to take-up or switch to ISV products.

- c) **Committed spend discounts:** All three hyperscalers use committed spend discounts, 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. 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.
- 10.2 As set out in Section 6, our evidence to date gives us reasonable grounds for suspecting that each of these practices individually may have a material influence on customer behaviour by reducing incentives to switch and multi-source. However, there is the potential for these practices to have an even greater impact when they combine with each other. Our provisional view is that these features are already harming public cloud infrastructure services customers today.
- 10.3 Looking ahead, we see a risk that the market becomes more concentrated compared to today, potentially around the market leaders' ecosystems. Where smaller cloud providers pose a weaker constraint on the ecosystems of the market leaders, this could dampen competition and lead to worse outcomes for customers in future.
- 10.4 Given the importance of cloud computing for digital services and the UK economy, our provisional view is that action should be considered now, both to improve outcomes today and mitigate future risks. Therefore, we are publishing a consultation on our proposal to make an MIR to the CMA alongside this report. Our consultation sets out the reasoning behind our provisional view that the legal threshold for an MIR is met, as well as our assessment of the proportionality of a market investigation. In making this decision to consult, we have taken into account the fact that market investigations can pose a significant burden on relevant industry players and the CMA. We welcome views on this preliminary position and we will issue our final decision alongside publication of our final report.
- 10.5 At the conclusion of a market investigation, the CMA group will consider whether any remedies are needed. If they are, they may impose a broad range of remedies, including market-opening measures, structural measures, or recommendations to Government or other regulatory bodies to change policy, legislation, or regulatory frameworks. We have identified here some potential interventions that could be implemented through such routes, but it would ultimately be up to the CMA group to assess their proportionality and consider their effectiveness alongside other suitable remedies they may identify, where appropriate.⁷⁸⁹

⁷⁸⁹ See Section 133A(1)(a) EA02 and the sections referenced therein including in particular section 134.

A1. Responding to this consultation

How to respond

- A1.1 Ofcom would like to receive views and comments on the issues raised in this document, by 5pm (UK time) on 17 May 2023.
- A1.2 You can download a response form from <https://www.ofcom.org.uk/consultations-and-statements/category-3/cloud-services-market-study>. You can return this by email or post to the address provided in the response form.
- A1.3 If your response is a large file, or has supporting charts, tables or other data, please email it to cloudreport@ofcom.org.uk, as an attachment in Microsoft Word format, together with the [cover sheet](#). This email address is for this consultation only, and may not be valid after 31 May 2023.
- A1.4 Responses may alternatively be posted to the address below, marked with the title of the consultation:
- Cloud services team
Ofcom
Riverside House
2A Southwark Bridge Road
London SE1 9HA
- A1.5 We welcome responses in formats other than print, for example an audio recording or a British Sign Language video. To respond in BSL:
- send us a recording of you signing your response. This should be no longer than 5 minutes. Suitable file formats are DVDs, wmv or QuickTime files; or
 - upload a video of you signing your response directly to YouTube (or another hosting site) and send us the link.
- A1.6 We will publish a transcript of any audio or video responses we receive (unless your response is confidential)
- A1.7 We do not need a paper copy of your response as well as an electronic version. We will acknowledge receipt of a response submitted to us by email.
- A1.8 You do not have to answer all the questions in the consultation if you do not have a view; a short response on just one point is fine. We also welcome joint responses.
- A1.9 It would be helpful if your response could include direct answers to the questions asked in the consultation document. The questions are listed at Annex 4. It would also help if you could explain why you hold your views, and what you think the effect of Ofcom's proposals would be.
- A1.10 If you want to discuss the issues and questions raised in this consultation, please contact Warwick Izzard on 020 7783 4127, or by email to warwick.izzard@ofcom.org.uk.

Confidentiality

- A1.11 Consultations are more effective if we publish the responses before the consultation period closes. In particular, this can help people and organisations with limited resources or familiarity with the issues to respond in a more informed way. So, in the interests of transparency and good regulatory practice, and because we believe it is important that everyone who is interested in an issue can see other respondents' views, we usually publish responses on [the Ofcom website](#) at regular intervals during and after the consultation period.
- A1.12 If you think your response should be kept confidential, please specify which part(s) this applies to, and explain why. Please send any confidential sections as a separate annex. If you want your name, address, other contact details or job title to remain confidential, please provide them only in the cover sheet, so that we don't have to edit your response.
- A1.13 If someone asks us to keep part or all of a response confidential, we will treat this request seriously and try to respect it. But sometimes we will need to publish all responses, including those that are marked as confidential, in order to meet legal obligations.
- A1.14 To fulfil our pre-disclosure duty, we may share a copy of your response with the relevant government department before we publish it on our website.
- A1.15 Please also note that copyright and all other intellectual property in responses will be assumed to be licensed to Ofcom to use. Ofcom's intellectual property rights are explained further in our [Terms of Use](#).

Next steps

- A1.16 Following this consultation period, Ofcom plans to publish a final report no later than 5 October 2023.
- A1.17 If you wish, you can [register to receive mail updates](#) alerting you to new Ofcom publications.

Ofcom's consultation processes

- A1.18 Ofcom aims to make responding to a consultation as easy as possible. For more information, please see our consultation principles in Annex 2.
- A1.19 If you have any comments or suggestions on how we manage our consultations, please email us at consult@ofcom.org.uk. We particularly welcome ideas on how Ofcom could more effectively seek the views of groups or individuals, such as small businesses and residential consumers, who are less likely to give their opinions through a formal consultation.
- A1.20 If you would like to discuss these issues, or Ofcom's consultation processes more generally, please contact the corporation secretary:

Corporation Secretary
Ofcom
Riverside House
2a Southwark Bridge Road
London SE1 9HA
Email: corporationsecretary@ofcom.org.uk

A2. Ofcom's consultation principles

Ofcom has seven principles that it follows for every public written consultation:

Before the consultation

- A2.1 Wherever possible, we will hold informal talks with people and organisations before announcing a big consultation, to find out whether we are thinking along the right lines. If we do not have enough time to do this, we will hold an open meeting to explain our proposals, shortly after announcing the consultation.

During the consultation

- A2.2 We will be clear about whom we are consulting, why, on what questions and for how long.
- A2.3 We will make the consultation document as short and simple as possible, with an overview of no more than two pages. We will try to make it as easy as possible for people to give us a written response.
- A2.4 We will consult for up to ten weeks, depending on the potential impact of our proposals.
- A2.5 A person within Ofcom will be in charge of making sure we follow our own guidelines and aim to reach the largest possible number of people and organisations who may be interested in the outcome of our decisions. Ofcom's Consultation Champion is the main person to contact if you have views on the way we run our consultations.
- A2.6 If we are not able to follow any of these seven principles, we will explain why.

After the consultation

- A2.7 We think it is important that everyone who is interested in an issue can see other people's views, so we usually publish the responses on our website at regular intervals during and after the consultation period. After the consultation we will make our decisions and publish a statement explaining what we are going to do, and why, showing how respondents' views helped to shape these decisions.

A3. Consultation coversheet

BASIC DETAILS

Consultation title:

To (Ofcom contact):

Name of respondent:

Representing (self or organisation/s):

Address (if not received by email):

CONFIDENTIALITY

Please tick below what part of your response you consider is confidential, giving your reasons why

Nothing

Name/contact details/job title

Whole response

Organisation

Part of the response

If there is no separate annex, which parts? _____

If you want part of your response, your name or your organisation not to be published, can Ofcom still publish a reference to the contents of your response (including, for any confidential parts, a general summary that does not disclose the specific information or enable you to be identified)?

DECLARATION

I confirm that the correspondence supplied with this cover sheet is a formal consultation response that Ofcom can publish. However, in supplying this response, I understand that Ofcom may need to publish all responses, including those which are marked as confidential, in order to meet legal obligations. If I have sent my response by email, Ofcom can disregard any standard e-mail text about not disclosing email contents and attachments.

Ofcom aims to publish responses at regular intervals during and after the consultation period. If your response is non-confidential (in whole or in part), and you would prefer us to publish your response only once the consultation has ended, please tick here.

Name

Signed (if hard copy)

A4. Consultation questions

Question 4.1 Do you agree with our assessment of how customers buy cloud infrastructure services and how cloud providers seek to acquire customers?

Question 4.2: Do you agree with our characterisation of the market outcomes in supply of cloud infrastructure services?

Question 5.1: Do you agree with our analysis of potential barriers to switching and multi-cloud? As part of this:

a) Please provide your views on the extent to which, and in what ways, egress fees are a barrier to switching and multi-cloud. Please also provide your views on the extent to which egress fees currently charged relate to the incremental cost of providing egress.

b) Please provide your views on whether specific business practices of cloud providers, particularly the hyperscalers, exacerbate technical barriers to switching and multi-cloud.

c) Please provide your views on how committed spend discounts are set and the impact these discounts have on the incentives of customers to multi-cloud.

Question 5.2: Do you agree with our analysis of potential barriers to entry and expansion?

Question 5.3: Do you agree with our analysis of the hyperscalers' relationship with ISVs? As part of this, please provide your views on whether our analysis of the hyperscalers relationship with ISVs applies to both larger and smaller ISVs.

Question 6.1 Do you agree with our assessment of how well competition is working in cloud infrastructure and what are the potential implications of a lack of competition?

Question 8.1 Do you agree that egress fees are an area of potential intervention? How might such an intervention be approached?

Question 8.2: Do you agree that interoperability and portability are areas of potential intervention? How might such an intervention be approached?

Question 8.3: Do you agree that committed spend discounts are an area of potential intervention? How might such an intervention be approached?

Question 8.4: Do you agree that transparency of billing is an area of potential intervention? How might such an intervention be approached?

Question 8.5: What, if any, potential unintended consequences do you anticipate might be associated with the interventions set out above, and how might they interact with each other if implemented?