Cloud services market study

Final report – annexes

Redacted [✂] for publication

Published 5 October 2023

Contents

Annexes

A1. Share of supply for cloud infrastructure services in the UK 3
A2. Profitability of the hyperscalers 14
A3. Prevalence of multi-cloud and switching 48
A4. Further detail on technical barriers to multi-cloud and switching 77
A5. Discounting analysis 99
A6. Terms of reference 110
A7. Glossary 111

A1. Share of supply for cloud infrastructure services in the UK

A1.1 This annex sets out our analysis of UK shares of supply for cloud infrastructure services (i.e. IaaS and PaaS), focusing on the shares of the hyperscalers (AWS, Microsoft Azure and Google Cloud Platform - GCP).

A1.2 In this annex we present our estimates of the following:

a) The revenue associated with cloud infrastructure services in the UK.
b) Shares of supply for cloud infrastructure services in the UK, based on revenues.[1]

A1.3 We also comment on which types of products represent most revenue for the hyperscalers in the UK.

A1.4 This annex updates analysis presented in Annex 5 of our interim report to include 2022 data and responds to stakeholder comments on our share of supply analysis.

Source of data

A1.5 Our analysis is based on responses to our statutory information requests under section 174 of the Enterprise Act 2002 and information purchased from third party data providers Synergy Research Group ('Synergy') and International Data Corporation ('IDC').

A1.6 We first explain these three sources of data, before describing our methodology for using them to estimate UK revenues and shares of supply for cloud infrastructure services.

Statutory information requests

A1.7 We asked 11 companies for information on their annual UK public cloud revenues between 2019 and 2022 and their global cloud revenues in 2021 and 2022. We approached these companies as they appeared to account for a significant proportion of UK revenues in the IaaS and PaaS segments according to the data we purchased from Synergy and IDC.

A1.8 We asked these 11 companies for revenue on each of their public cloud services, with services categorised into i) product categories as per their website (e.g. compute, storage, etc) and ii) IaaS, PaaS, SaaS segments and sub-segments.[2]

A1.9 We requested UK data based on billings (i.e. if a customer is billed in the UK, we asked companies to include all cloud revenues associated with that customer).

A1.10 The companies we approached could not always provide the data in the requested format. For example, not all companies provided revenues by service, product category and

[1] Estimates of shares of supply from companies like Gartner, IDC and Synergy are generally based on revenue.
[2] For IAAS, we asked providers to categorise their services into compute, storage, networking and other. For PaaS services, we asked providers to categorise their services into application development platforms, containers, functions (serverless), data management, data analytics/AI/ML, IOT and other. We did not request any SaaS sub-segment data.

IaaS/PaaS categorisation. Similarly, while all companies provided estimates of UK revenues, in some cases these did not align to what customers were billed or how revenue was recognised in statutory accounts.[3]

A1.11 The table below summarises the revenue data provided by the 11 companies we sent information requests to.[4]

Table A1.1: Summary of revenue information received in response to information requests

Company	Revenue for years ending	Revenue provided by:		
		Service	Product category	IaaS/PaaS
Amazon (AWS)	Dec	Some	✓	×
Microsoft (Azure)	Dec	✓	✓	✓
Google (GCP)	Dec	Some	✓	✓
Oracle	May	Some	✓	✓
IBM	Dec	×	✓	✓
OVH	Dec	Some	✓	✓
Snowflake	Jan	✓	✓	✓
MongoDB	Jan	✓	✓	✓
Salesforce	Jan	✓	✓	✓
VMware	Dec	Some	×	✓
Atlassian	June	Some	✓	✓

Note: under the column headed 'service', we use 'some' to denote that revenue was provided but not always split out into individual services.

A1.12 The table below sets out some of the main limitations of the data we received in response to our information requests and the adjustments we have made, if any, for the purposes of our shares of supply analysis. Most of these limitations are relatively minor and do not significantly affect our share of supply estimates. Limitations around the attribution of revenue between IaaS and PaaS services could affect our share of supply estimates for IaaS and PaaS, but are less likely to affect our estimates for IaaS and PaaS combined. We do not consider that any of these limitations would affect our overall findings that AWS and Microsoft have the largest shares of supply for IaaS and PaaS in the UK.

[3] Some data received recognised all revenue associated with a contract at the date of purchase, rather than over time as services are delivered. Other data did not include free credit discounts and therefore differed from the amount billed to customers.

[4] Responses by provider: [✂]

Table A1.2: Data from information requests: limitations and adjustments

Limitation	Adjustments to the data
Some providers were unable to separate private and public cloud revenues	We have not attempted to make any adjustment for this. Private cloud revenues are a small proportion of cloud revenues for providers who were able to separate the two. Qualitative responses from other providers also indicated private cloud revenues are small relative to public cloud.
Attribution of revenues to IaaS, PaaS and SaaS categories	We generally relied on cloud provider categorisations provided in responses. Although we provided definitions of these categories in our information requests, we recognise that respondents could interpret these classifications differently, which could affect our shares of supply analysis. We have categorised revenues ourselves where providers were unable to do so. Where we have needed to do this, we have attributed compute, storage and networking services to IaaS and remaining services to PaaS.
Treatment of IaaS revenues triggered by PaaS usage.	Where a customer uses a service, Service A (e.g. a PaaS service) that is built on an underlying service, Service B (e.g. an IaaS service), providers have different charging approaches. For example, AWS charges the customer for the fees incurred with the usage of Service A and does not additionally charge the customer for the underlying usage of Service B. Responses indicated that some providers attribute IaaS fees triggered by PaaS usage to the PaaS service (for example AWS and Google), while others attribute such IaaS fees to IaaS services (for example Microsoft). We asked hyperscalers how much compute and storage revenue was driven by database and analytics use (which are two of the larger PaaS services). AWS said these services account for [✂] of total compute and storage revenues, [✂].[5] Microsoft and Google were not able to quantify the amount of compute and storage revenues associated with usage by these services.[6] We are unable to adjust the data to ensure IaaS revenues are treated in the same way across providers (whether for data provided in response to information requests or any other data sources we are using). This means that IaaS and PaaS revenues for different providers may not be comparable in some cases, which could affect our share of supply analysis. Microsoft commented that this different approach to the treatment of certain IaaS fees could mean

[5] AWS response dated 17 February 2023 to our follow-up email dated 28 January 2023 concerning the s.174 notice dated 24 October 2022, question 2.
[6] Microsoft response dated 10 February 2023 to question 4 of our follow up email dated 27 January 2023 concerning the s.174 notice dated 21 October 2022, Part B question 11. Google response dated 10 February 2023 to our follow-up email dated 27 January 2023 concerning the s.174 notice dated 26 October 2022, question 2.

Limitation	Adjustments to the data
	its IaaS share of supply is 'over-estimated'.[7] As we understand Microsoft records all IaaS fees triggered by PaaS usage against IaaS services and AWS and Google do not, we agree that IaaS revenues for Microsoft could appear relatively high, and PaaS revenues relatively low compared to AWS and Google. However, based on our understanding of the market, we would not expect these different approaches to have a material impact on our shares of supply analysis for IaaS and PaaS. UK revenues and shares of supply for IaaS and PaaS combined will be unaffected by differences in choices on how to record revenue between IaaS and PaaS.
Some companies provided data on a financial year rather than calendar year basis	We converted financial year data to calendar years pro-rata. The hyperscalers all provided data in calendar years.
Most revenue data was provided in non-GBP currencies (US dollars, Euro)	We converted non-GBP currencies into GBP using ONS average monthly exchange rates for the relevant time period.
Incomplete data for certain years or geographies	• One provider could only provide aggregate cloud revenue data for 2019. We split this into categories using this provider's 2020 proportions. • One provider could not provide UK revenue data for its 2022 financial year. We estimated UK IaaS and PaaS revenues for FY2022 using its global revenue growth rates for IaaS and PaaS in this year. • One provider did not provide data for its 2019 financial year, which meant our estimated calendar year revenues did not cover all of 2019. We have made no adjustment for this as it will not materially affect our shares of supply analysis.
Mismatch between total revenues and product category revenues	[✂] pulled data from two separate internal systems to respond to our information requests.[8] These systems are used for different purposes and therefore adopt a slightly different definition of revenue. The first system is used as a basis for preparing financial statements and is [✂] preferred way of providing financial data where possible, but only includes a limited breakdown of total cloud revenues. The second system included revenues by product category, but total revenues did not reconcile to the first system. The difference in total revenues between these systems in 2020 to 2021 was small, so in our interim report we used data from the second system to estimate shares of supply. However, in 2022 the difference was more significant. As a result, we have scaled product category revenues so that total revenues used in our analysis match those from the first system. We have applied this change to

[7] Microsoft response to the interim report, page 32, paragraph 3.
[8] [✂]

Limitation	Adjustments to the data
	each year 2020 to 2022. These changes impact the provider's share by less than 0.5 percentage points so do not have a material impact on our shares of supply for 2020 to 2021.

Synergy and IDC UK data

A1.13 It was not practical to send a statutory information request to all companies providing cloud services in the UK. To assess overall UK revenues for IaaS and PaaS we therefore purchased data from third-parties. We bought data from two providers, Synergy and IDC to understand the extent to which our analysis may be sensitive to different estimates of total UK revenues.

A1.14 We also purchased updated UK data after our interim report. Synergy and IDC included restatements of market size and shares of supply for previous years. For IDC these years are 2019 – 2021 and for Synergy they are 2020 - 2021. We have taken account of these restatements, though they do not have a material impact on our shares of supply analysis.

Synergy

A1.15 We purchased UK revenue data on IaaS and PaaS from Synergy for the calendar years 2018 to 2022.[9]

A1.16 Synergy's IaaS data is broken down into storage, compute, networking, and other.

A1.17 Synergy's PaaS data is broken down into analytics, databases, IoT (internet of things), and other.

A1.18 In each category, Synergy estimates UK revenue (in USD) and shares of supply associated with a number of companies (around 40 in IaaS and 20-30 in PaaS). We do not have a mapping of individual cloud services to these categories, though Synergy provided us with some examples of hyperscaler services included in these categories.

IDC

A1.19 We bought UK revenue data on the following 'primary markets' from IDC: IaaS, Application Development and Deployment (PaaS), Systems Infrastructure Software (SaaS), and Applications (SaaS).[10] This data covered the calendar years 2017 to 2022.

A1.20 IDC's IaaS data is broken down into three secondary markets: storage, compute, and networking.

[9] Synergy Research Group, 4Q 2021 Cloud Infrastructure Services United Kingdom Market Share Report, March 2022 and Synergy Research Group, 4Q 2022 Cloud Infrastructure Services United Kingdom Market Share Report, March 2023. Our analysis focuses on public cloud infrastructure and only includes public IaaS and public PaaS services. We have not included managed private cloud services, which is also included in Synergy's Cloud Infrastructure Services dataset. We converted Synergy's data from US dollars to pound sterling using average monthly exchange rates from ONS for the relevant period.

[10] IDC, Public Cloud Services Tracker 2021 H2 (published April 2022) and IDC, Public Cloud Services Tracker 2022 H2 (published April 2023). We converted IDC's data from US dollars to pound sterling using average monthly exchange rates from ONS for the relevant period.

A1.21 IDC's PaaS data is broken down into seven secondary markets: analytics and business intelligence, AI platforms, data management, integration and orchestration, application development, software quality and life cycle, and application platforms.

A1.22 We did not purchase any additional granularity on the secondary markets for SaaS - Systems Infrastructure Software or SaaS - Applications.

A1.23 In each category, IDC estimates UK revenue (in USD) and shares associated with around 30 companies in IaaS and over 200 in PaaS. We do not have a mapping of individual cloud services to these primary or secondary markets, though IDC provided us with copy of its taxonomy documents which describe these markets and give examples of products in each category.[11]

UK revenues for IaaS and PaaS

Methodology

A1.24 To estimate total UK revenues for IaaS and PaaS we did the following:

a) We assumed data provided in response to our statutory information requests represented the best available estimate of UK revenues on IaaS and PaaS for the 11 companies we sent these to. We adjusted this data as set out in Table A1.2 above. Total UK revenue from these 11 companies in 2022, across both IaaS and PaaS, was around £5.9bn.
b) We estimated the remainder of UK revenues associated with IaaS and PaaS using Synergy and IDC data (excluding their revenue estimates for the 11 providers we had revenue data for from our information requests).

A1.25 This approach gave us two estimates for the total UK revenues associated with IaaS and PaaS: one based on a combination of our information request data and Synergy data, and another based on a combination of our information request data and IDC data.

A1.26 When comparing these two estimates, total revenues associated with IaaS in the UK are similar, though there is a larger difference between the estimates of revenue associated with PaaS. This could be due to a few factors including:

a) potential differences in definitions of IaaS, PaaS and SaaS, and views into which category individual services fit,[12]
b) differences in the number of companies tracked, and
c) potential differences in source data, modelling methodology, and approach to estimating UK revenues.

A1.27 In the table below we present the average of these UK revenue estimates. We think it is reasonable to present the average as, for PaaS in particular, [✂].[13] For publication, we report these averages within ranges of £500m.

[11] IDC Worldwide Software Taxonomy, 2022, IDC Worldwide Software Taxonomy, 2023 and IDC Worldwide IT Cloud Services Taxonomy, 2019.
[12] We recognise that the categorisation of services into IaaS/PaaS/SaaS could affect estimates of shares of supply, and that categorisation is not always easy as some services will be on the boundary of more than one category, or there could be different views on how to appropriately categorise services.
[13] [✂].

UK revenues for IaaS and PaaS

A1.28 We estimate that in 2022, cloud infrastructure services generated revenues of [✂] [£7.0bn to £7.5bn]. Between 2019 and 2022, UK revenues for IaaS and PaaS combined grew by 35% - 40% per year.

Table A1.3: UK IaaS and PaaS revenues, 2019 – 2022 (£bn)

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

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

A1.29 Within IaaS, the main product categories are compute, storage and networking. In 2022 we estimate that these represented around 65%, 20% and 8% of UK IaaS revenues respectively.[14] Since 2019 the revenues in each category have grown. We estimate that compute revenues have increased the most in absolute terms, though storage and networking have grown faster in percentage terms.

A1.30 The product categories associated with PaaS vary between those provided in response to our information requests, and those used by Synergy and IDC data, making it difficult to estimate UK revenues for individual categories. However, based on our analysis of information provided in response to our information requests, services associated with data management and analytics appear to generate the most revenues, representing over [✂] of UK PaaS revenue for hyperscalers in aggregate in 2022.

UK shares of supply for IaaS and PaaS

Methodology

A1.31 To estimate UK shares of supply for IaaS and PaaS we did the following:

a) We estimated UK shares of supply for the 11 providers we sent information requests to, based on the revenue information they provided. In the tables below we only show individual shares of supply for AWS, Microsoft (for Azure) and Google (for GCP), though we also comment on the shares of supply for some of the larger providers outside of the hyperscalers.
b) We estimated shares of supply by taking cloud service revenues for these companies (as provided in response to our information requests) and dividing by our total UK revenue estimates shown in Table A1.3.

[14] These estimates do not sum to 100% as not all respondents to our information requests provided a breakdown of IaaS revenues by product category and the Synergy data (which we used to inform our UK revenue estimates) includes an 'IaaS other' category.

c) For publication, our shares of supply estimates have been placed within ranges of ten percentage points (five percentage points where our estimate is below 10%).[15]

UK shares of supply for IaaS and PaaS

IaaS

A1.32 The table below shows our estimated shares of supply for UK IaaS.

Table A1.4: UK IaaS shares of supply, 2019 – 2022

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

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

A1.33 We estimate that in 2022 AWS and Microsoft represented approximately [✂] [80% to 90%] of UK IaaS revenues, a percentage that has increased by around [✂] [0 to 5] percentage points since 2019. Within this, we estimate that Microsoft's share of UK IaaS revenues has grown while AWS's share has reduced slightly. As noted in Table A1.2, AWS and Google allocate IaaS fees triggered by PaaS to PaaS services, which is a different approach to Microsoft. As a result, their UK IaaS revenues could be relatively lower than Microsoft's.

A1.34 While Google's IaaS revenues in the UK have grown quickly since 2019, in 2022 we estimate it represented [✂] [0% to 5%] of UK IaaS revenues, significantly behind AWS and Microsoft.

A1.35 Within the 'other' category, IBM and Oracle appear to have the next largest shares of supply, though these are both [✂] [0% to 5%]. While UK IaaS revenues for smaller cloud providers have grown since 2019, [✂].

A1.36 AWS said that its "infrastructure" share of supply had fallen significantly between 2016 and 2021.[16] It referenced Gartner analysis indicating a fall in AWS' global IaaS share from 44.2% in 2016 to 38.9% in 2021, but did not provide additional data on its UK IaaS share.[17] Although the data we requested as part of our information requests and obtained from IDC and Synergy only allows us to estimate UK shares of supply between 2019 and 2022, our

[15] As noted above, the restated data from Synergy and IDC for the years 2019, 2020 and 2021 did not have a material impact on our shares of supply analysis. However, there are two instances where our refined share of supply estimate moved the result from one five-percentage-point band to another (e.g. from 10-20% to 5-10%). Where this is the case, we have not changed the five-percentage-point band reported in the interim report to preserve the confidentiality of our point estimates. This does not affect any of our findings presented in this annex.

[16] AWS response to the interim report, page 3, paragraph 5.

[17] Gartner website. Gartner Says Worldwide IaaS Public Cloud Services Market Grew 31 Percent in 2016 [accessed 29 September 2023] and Gartner Says Worldwide IaaS Public Cloud Services Market Grew 41.4% in 2021 [accessed 29 September 2023].

analysis suggests AWS' share of UK IaaS revenues fell slightly between 2019 and 2022, although as of 2022 it remained the largest provider.

PaaS

A1.37 PaaS includes many diverse types of services. There are also more companies and ISVs providing PaaS compared to IaaS, and many of these 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. For example, in 2022, only the hyperscalers and IBM were active in each of IDC's seven 'secondary markets'[18] while Oracle operated across most of these categories. Most other companies tracked by IDC were only active in one or two secondary markets.

A1.38 We think there is value in showing PaaS shares of supply for the hyperscalers to compare their positions to IaaS – particularly as they are the key providers who offer services across all application segments within PaaS. However, we acknowledge looking at PaaS overall may mask differences in shares between sub-categories of PaaS, as some sub-categories offer more provider alternatives than others.

A1.39 The table below shows our estimated shares of supply for UK PaaS.

Table A1.5: UK PaaS shares of supply, 2019 – 2022

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

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

A1.40 We estimate that AWS and Microsoft represented [✂] [50% to 60%] of UK PaaS revenues in 2022 – a lower share than for IaaS. Within this, we estimate that AWS and Microsoft's share of UK PaaS revenues has grown, though Microsoft's share has grown slightly faster. As noted in Table A1.2, we understand that Microsoft allocates IaaS revenues triggered by PaaS to IaaS services, which is a different approach compared to AWS and Google. As a result, Microsoft's PaaS revenues could be relatively lower.

A1.41 We estimate that Google's share is closer to that of AWS and Microsoft in PaaS than in IaaS, with a [✂] [10% to 20%] share of UK PaaS revenues in 2022 and its share increasing since 2019. [✂].

A1.42 The "other" category is very broad and includes many companies with low shares of supply. The share of supply of the "other" category fell in each year between 2019 and 2022, and by around [✂] [10% to 20%] points over the period as a whole. Oracle, MongoDB and IBM have some of the larger shares of supply of companies in the 'other' category; we estimate

[18] As noted above, these are analytics and business intelligence, AI platforms, data management, integration and orchestration, application development, software quality and life cycle, and application platforms.

they represented around [✂] [0% to 5%], [✂] [0% to 5%] and [✂] [0% to 5%] respectively of UK PaaS revenues in 2022. Over this period, UK PaaS revenue growth for smaller cloud providers generally lagged behind hyperscaler revenue growth.

A1.43 We estimate that ISVs collectively accounted for [✂] [30% to 40%] of UK PaaS revenues in 2022.

IaaS and PaaS combined

A1.44 Table A1.6 shows our estimated shares of supply for UK IaaS and PaaS combined, drawing on the information presented above. We estimate that in 2022 AWS and Microsoft between them had [✂] [70% to 80%] share of UK IaaS and PaaS revenues, with Google significantly lower on [✂] [5% to 10%]. Overall, we estimate that AWS, Microsoft and Google accounted for [✂] [70% to 80%] of UK IaaS and PaaS revenues in 2022.

A1.45 The "other" category includes many companies with low shares of supply. IBM and Oracle have some of the larger shares of supply of companies in the 'other' category; we estimate they both represented around [✂] [0% to 5%] of UK IaaS and PaaS (combined) revenues in 2022 and neither have grown significantly since 2019.

A1.46 Between 2019 and 2022, we estimate that AWS' UK share of supply fell marginally, while the UK share for Microsoft and Google grew. Google's share experienced stronger growth in relative terms, although from a lower revenue base.

Table A1.6: UK shares of supply for IaaS and PaaS combined, 2019 – 2022

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

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

A1.47 AWS said that, according to IDC data, a number of smaller providers of UK cloud infrastructure services significantly increased revenues over the period 2017 to 2022.[19] Our analysis of UK revenues for IaaS and PaaS indicates that, while smaller cloud providers increased UK revenues between 2019 and 2022, they lost share as their revenue growth was slower than that of hyperscalers. Indeed, since 2019 we estimate 80-90% of UK revenue growth for IaaS and PaaS combined came from AWS, Microsoft and Google.[20]

[19] AWS response to the interim report, page 3, paragraph 5. AWS references IDC's Worldwide Semiannual Public Cloud Services Tracker, H2 2022.
[20] Based on our analysis of data provided in response to our information requests and data from Synergy and IDC.

Major product categories for the hyperscalers

A1.48 Hyperscalers provide hundreds of cloud services, but based on responses to our statutory information requests, we estimate that their UK revenue is concentrated in five types of services.

A1.49 As stated above, the hyperscalers provided us with revenue data and, at our request, attributed this to product categories used on their websites and to IaaS and PaaS categories that we provided.[21] Based on the information provided, we identified revenues associated with five broad categories: compute, storage, networking, data management and analytics.[22]

A1.50 Figure A1.7 shows the proportion of UK IaaS and PaaS revenues that we estimate was associated with these five categories in 2022.[23] Consistent with the analysis of 2021 data presented in the interim report, it indicates that:

a) Compute is typically the largest product category.
b) Data management services are typically the second or third largest product category.
c) Storage is relatively important for Microsoft.
d) [✂] for GCP.

Figure A1.7 Proportion of hyperscalers' UK IaaS and PaaS revenue from particular product categories, 2022

[✂]

Source: Ofcom analysis of data provided in response to our information requests. Note that the chart does not include an 'other' category so does not sum to 100%.

A1.51 We estimate that compute was the slowest growing category in the UK across the hyperscalers between 2019 and 2022 in percentage terms but it had the largest absolute growth. The relative share of UK revenue from other categories, particularly networking and analytics, increased during this period.

A1.52 We have less information on smaller cloud providers, but based on the data provided, we estimate that they also appear to generate most UK revenues from these five product categories.

[21] Microsoft noted that some services could reasonably fit into multiple product categories. Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, question 1.
[22] Our analysis was based on the following responses: AWS [✂]; Google [✂]; and Microsoft [✂].
[23] Some of the larger types of product categories not included in the chart are hybrid cloud management (e.g. Google Anthos, Azure Stack), Kubernetes (e.g. Google Kubernetes Engine) and management and governance services.

A2. Profitability of the hyperscalers

Introduction

A2.1 This annex updates the profitability analysis set out in Annex 6 of our interim report to include the latest available data and to respond to stakeholder comments on our profitability analysis.

A2.2 As set out in the interim report, this annex presents our analysis of the profitability of the hyperscalers. We are particularly focused on the profitability of their cloud infrastructure businesses: AWS for Amazon, Azure for Microsoft and Google Cloud Platform for Google.

A2.3 Profitability can be one indicator of how well competition is working and we consider it alongside other indicators in Section 8. We set out our analysis of the following in this annex:

a) We compare hyperscaler operating profits and margins to those of other cloud providers. If hyperscaler profits are significantly higher than other providers, or if other providers are unable to make a profit, this could be consistent with the existence of economies of scale and scope and indicate that it is difficult for competitors to sustainably enter the cloud services market. Operating margins on their own do not indicate whether returns are higher than might be expected in a market that is working well. We therefore also compare returns to the cost of capital, as described below.
b) We assess whether the hyperscalers' cloud businesses have generated returns persistently above their weighted average cost of capital (WACC). Evidence of this may reflect limitations to the competitive process, where barriers to effective competition may limit the competitive constraints larger firms face. Returns above the cost of capital in any particular period are not necessarily indicative of a competition problem. However, where returns are persistently high they can, in combination with other evidence, suggest competition is not working well.

A2.4 In the interim report, we set out our analysis of profitability in Annex 6 and considered what it might tell us about the effectiveness of competition in Section 6. Most stakeholder comments focused on the relevance of our analysis to the market study and the interpretation of our findings rather than our analysis itself or how it was prepared. Therefore, in the sub-sections below we do two things. First, we update our profitability analysis with the latest data and respond to any comments on how we carried out our analysis. Updating our analysis does not materially affect any of our findings from the interim report. Second, we respond to stakeholder comments on the relevance of our analysis to the market study and the interpretations of our findings.

A2.5 This annex is split into the following sub-sections:

a) **Source of data and scope of analysis.** This is largely the same as presented in the interim report.
b) **Hyperscaler cloud profitability**: We assess hyperscalers' cloud revenues and operating profits and compare these to other cloud infrastructure providers. This sub-section has been updated to include the most recent annual financial statements from Microsoft, Oracle and Alibaba, and quarterly results from all cloud providers. It also reflects

responses to information requests sent to cloud providers after our interim report. Our findings are similar to those presented in the interim report.

c) **Assessment of ROCE relative to WACC**: We estimate the return on capital employed (ROCE) earned by AWS and Microsoft on their cloud investments and compare this to the WACC. This sub-section has been updated to include the most recent annual financial statements from Microsoft, quarterly results from AWS and more recent responses to our information requests. Our findings are similar to those presented in the interim report. We respond to stakeholder comments on our ROCE analysis in this sub-section.

d) **Relevance and interpretation of our findings.** We respond to stakeholder comments on the relevance and interpretation of our findings, including whether returns above WACC could represent rewards for successful risk taking.

Source of data and scope of analysis

Source of data

A2.6 We carried out 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. For example, UK customers can and do use cloud infrastructure services hosted in data centres around the world, and they purchase products which are generally available across the hyperscalers' global customer base. This is reflected in the fact that much of the financial data we have analysed (particularly cost and profit information) is only available at the global level. We did request information on UK costs and profits associated with cloud to supplement our global analysis, but companies were generally unable to provide this.[24]

A2.7 Our analysis is based on global information from published financial statements and information provided by cloud providers in response to our statutory information requests.

A2.8 Cloud providers' reporting of their cloud businesses varies in granularity and specificity. The focus of our market study is cloud infrastructure services. In many cases, providers' financial reporting is less specific than this and the cloud infrastructure business is grouped together with other businesses for reporting purposes.

A2.9 Table A2.1 below lists the hyperscalers and other global cloud infrastructure providers mentioned in this annex and sets out for each of them our understanding of the name of their cloud infrastructure services business, within which operating segment this is reported in their financial statements, and whether that segment also includes non-cloud activities.

[24] In recent market studies, the CMA has also assessed profitability on a global basis. For example, in its 2020 Online platforms and digital advertising market study, the CMA considered global profits for Google and Facebook, and in its 2022 Mobile ecosystems market study, the CMA considered global profits for Apple and Google.

Table A2.1: Summary of cloud providers' public financial data reporting

Company	Cloud infrastructure services business	Reported segment featuring cloud infrastructure services business	Does reported segment include non-cloud infrastructure activities?
Alphabet	Google Cloud Platform (GCP)	Google Cloud	Yes – Google Cloud also includes Google Workspace's communication and collaboration tools, including apps such as Gmail, Drive and Meet
Amazon	AWS	AWS	No
Microsoft	Azure	Intelligent Cloud	Yes[25] – Intelligent Cloud also includes non-cloud server products and enterprise support and consulting services
IBM	IBM public cloud[26]	N/A (IBM previously reported revenue for 'hybrid cloud', which appears to span all of IBM's reported segments)	N/A
Oracle	Oracle Cloud Infrastructure	Cloud and License	Yes – Cloud and License also includes licensing and support activities for on-premises products like Oracle Database[27]
OVHcloud	Public Cloud	Public Cloud	No
Alibaba[28]	Alibaba Cloud	Cloud	Yes – some activities within the reported Cloud segment do not relate to infrastructure services[29]
DigitalOcean	DigitalOcean (whole business)	DigitalOcean (whole business)	No

Source: Ofcom analysis of cloud providers' published financial statements.

A2.10 Table A2.1 shows that for Alphabet, Microsoft, IBM, Oracle and Alibaba, the operating segments which include cloud infrastructure services also include some other cloud and/or non-cloud services. Microsoft Intelligent Cloud and Oracle Cloud and License both include

[25] Microsoft also reports revenue and gross profit margins for Microsoft Cloud, which is a collection of cloud-focused services from across Microsoft's reported operating segments (including Azure and other cloud services such as Office 365 Commercial, the commercial portion of LinkedIn, and Dynamics 365).
[26] IBM told us (in response to our information request) that it has a public cloud business within its hybrid cloud portfolio, which combines IaaS and PaaS services from its reported Infrastructure and Software segments.
[27] Oracle reports revenue for 'Infrastructure cloud services and license support' and 'cloud services', both of which, we understand, capture cloud infrastructure services, but also include some other services.
[28] Alibaba Cloud often appears as one of the largest cloud providers globally given its position in China. For example, IDC ranked it fourth globally in the provision of Foundational Cloud Services (which includes the IaaS, PaaS and SaaS – System Infrastructure Software market segments) in 2022.
[29] Alibaba Cloud specialises in the provision of cloud infrastructure and platform services. The wider Cloud segment also includes DingTalk, which is a communications and collaboration platform similar to Microsoft Teams.

services targeted at on-premises computing as well as cloud computing. This means that the financial performance of Microsoft Intelligent Cloud and Oracle Cloud and License may not reflect the financial performance of cloud infrastructure services.

A2.11 To obtain financial data specifically related to cloud infrastructure activities, we sent statutory information requests to some cloud providers asking for information on revenues, profit and capital employed. Table A2.2 summarises the information we asked for and what cloud providers were able to provide. While we were able to gather some additional data on cloud infrastructure services in relation to revenues and profit, providers could not provide capital employed information (beyond that reported in publicly available financial statements).

Table A2.2: Summary of financial data submitted by cloud providers in response to our information requests

Company	Data provided (for cloud infrastructure services business)			
	Revenue	Profit (income statement)	Capital employed (balance sheet)	Notes
Google	✓	Some	×	Google provided gross profits and operating profits for Google Cloud but only gross profits for Google Cloud Platform.[30]
AWS	✓	✓	×[31]	AWS provided a breakdown of AWS costs in more detail than that available publicly.[32]
Microsoft	✓	Some	×	Microsoft provided some profit information for Azure, but this did not include an allocation of all operating costs.[33]

[30] Google response dated 16 December 2022 to the s.174 notice dated 26 October 2022, Part B question 22; Google response dated 13 January 2022 to our follow-up email dated 21 December 2022 concerning the s.174 notice dated 26 October 2022, Part B question 9; Google response dated 26 May 2023 to our follow-up email dated 10 May 2023 concerning the s.174 notice dated 26 October 2022, Part B question 9.

[31] We note that Amazon does publish some asset information for AWS in its financial statements, as discussed further below.

[32] AWS response dated 13 January 2023 to question 5 of our follow-up email dated 16 December 2022 concerning the s.174 notice dated 24 October 2022, Part B question 22 (Annex Q5.1); AWS response dated 26 May 2023 to our follow-up email dated 5 May 2023 concerning the s.174 notice dated 24 October 2022, Part B question 9.

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

Company	Data provided (for cloud infrastructure services business)			
	Revenue	Profit (income statement)	Capital employed (balance sheet)	Notes
IBM	✓	Some	×	IBM commented on the operating profit for its public cloud business and provided only high-level estimates.[34]
Oracle	✓	×	×	Oracle said it could not provide profit information on cloud beyond its public reporting for the Cloud and License segment.[35]

Source: Ofcom analysis of cloud providers' responses to our statutory information requests. We did not request information from Alibaba or DigitalOcean. We did request revenue information from OVHcloud but not profit or capital employed.

Hyperscaler businesses we focus on for profitability analysis

A2.12 Given that reported data is not always specific to cloud infrastructure activities, and we have only been able to obtain limited additional data, our profitability analysis focuses on the following businesses of the hyperscalers:

a) **Amazon**: Our analysis references **AWS** as it provides cloud infrastructure services. The information available is largely sufficient for our analysis.

b) **Microsoft**: Our analysis primarily references **Azure** as it provides cloud infrastructure services. However, as the information available for Azure is limited, we also consider **Microsoft Cloud**, as Microsoft publishes more financial information on Microsoft Cloud compared to Azure. As well as Azure, Microsoft Cloud includes other cloud services like Office 365 Commercial, the commercial portion of LinkedIn, and Dynamics 365. While Azure is a part of Microsoft Cloud – in the year to June 2023 it represented around [✂] of Microsoft Cloud revenue[36] – Microsoft Cloud's financial performance reflects the performance of all Microsoft's cloud activities. Microsoft Cloud is not directly comparable with AWS for this reason.

c) **Google**: Our analysis references **Google Cloud**. While Google Cloud Platform is the part of Google Cloud focused on cloud infrastructure, we only have limited profitability data for this business. Google Cloud also includes Workspace, which incorporates Google's consumer and enterprise SaaS like Gmail and Google Docs (so Google Cloud could be comparable to Microsoft Cloud in terms of the suite of products it captures). Google Cloud is broader than 'cloud infrastructure services', and not directly comparable with

[34] 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; IBM response dated 4 April 2023 to the s.174 notice dated 25 October 2022, Part B question 9; IBM response dated 25 July 2023 to our follow-up email dated 5 May 2023 concerning the s.174 notice dated 25 October 2022, Part B question 9.

[35] Oracle response dated 16 December 2022 to the s.174 notice dated 31 October 2022, Part B question 20; Oracle response dated 6 January 2023 to our follow-up email dated 22 December 2022 concerning the s.174 notice dated 31 October 2022, Part B question 9 and Question 20; Oracle response dated 2 June 2023 to our follow-up email dated 10 May 2023 concerning the s.174 notice dated 31 October 2022, Part B question 9.

[36] Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22); Microsoft 10-K for the year ending June 2023, page 40.

AWS or Azure. However, we think it gives a reasonable idea of Google's financial performance in cloud infrastructure for the purpose of this market study because [✂].[37] We consider that Workspace is likely to be a smaller part of Google Cloud than Office 365 is to Microsoft Cloud (i.e. Google Cloud will better reflect the financial performance of cloud infrastructure activities than Microsoft Cloud).

Our approach to currency conversion for financial data

A2.13 Of the companies featured in Table A2.1, all of them report financial information in US dollars except Alibaba (Chinese Yuan) and OVHcloud (Euros). When calculating relative measures like revenue growth, profit margins and ROCE we use the reported currency. Where we present absolute monetary values, we have generally converted financial data into pound sterling using average exchange rates for the relevant period.[38]

Hyperscaler cloud profitability

A2.14 In this sub-section we set out the hyperscalers' cloud revenues and profits and compare these to other cloud providers. We consider the following:

a) relative contribution of cloud to hyperscalers' overall businesses
b) trends in hyperscaler cloud revenues
c) trends in hyperscaler cloud profit margins.

The relative contribution of cloud to hyperscalers' overall businesses

A2.15 Table A2.3 below shows the contribution of the hyperscaler cloud businesses to revenue and operating profit of the company as a whole (e.g. the percentage of Amazon's overall revenue and operating profit represented by AWS). For Microsoft, this table has been updated from the interim report for its results for the year ending June 2023.

Table A2.3: Cloud share of global revenue and operating profit (latest financial year)

Company and cloud business	Revenue share (%)	Operating profit share (%)
Amazon		
AWS	16%	186%
Alphabet		
Google Cloud	9%	-4%
Microsoft		

[37] Google response dated 31 March 2023 to the s.174 notice dated 26 October 2022, Part B question 4 (Annex 2).
[38] We specifically used the ONS Average Sterling exchange rate: US Dollar and ONS Average Sterling exchange rate: Euro time series for currency conversion where data is reported in US Dollars or Euros. For Alibaba's financial data, which is reported in Chinese Yuan, there is no ONS exchange rate data available, so we calculated financial year average exchange rates based on daily data from the Bank of England Database, and we used these calculated rates to convert data into pound sterling.

Company and cloud business	Revenue share (%)	Operating profit share (%)
Azure	[✂]	n/a
Microsoft Cloud	53%	n/a

Source: Ofcom analysis of cloud providers' published financial statements for financial years ending December 2022 (Amazon and Alphabet) and June 2023 (Microsoft) and information provided by Microsoft in response to our information requests.[39]

A2.16 While AWS generates a relatively small proportion of Amazon's revenues (16% in 2022), it is the main source of operating profit, due to lower (and sometimes negative) operating profit margins[40] across Amazon's other reported segments.[41] In contrast, Microsoft Cloud represented over half of Microsoft's revenue in its 2023 financial year. Microsoft does not publicly report operating profits for Azure or Microsoft Cloud. Microsoft's responses to our information requests indicate that Azure's share of total Microsoft revenue was [✂].[42] Google Cloud is a relatively small yet growing contributor to Alphabet's revenue. It has reported operating losses in its annual results to date (its latest results capture the financial year ending December 2022), although it made a small operating profit for the first time in the quarter ended March 2023, and a further operating profit in the quarter ended June 2023.[43]

Trends in hyperscaler cloud revenues

A2.17 Figure A2.4 below shows publicly reported global revenues associated with cloud activities for the three hyperscalers plus four global competitors (Alibaba, Oracle, DigitalOcean and OVHcloud), taken from their latest reported financial year. We comment on revenues to provide context for the profitability analysis in the next sub-section. This chart has been updated from the interim report for more recent annual results from Microsoft, Oracle and Alibaba.

[39] Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).
[40] Operating profit (EBIT) margin is calculated as operating profit divided by revenue.
[41] Amazon's North America and International segments largely consist of revenues from retail sales of consumer products. AWS has represented an increasing share of overall Amazon operating profits in recent years – 59% in 2020, 74% in 2021 and 186% in 2022, driven by higher profits for AWS and reduced profits in Amazon's retail operations in this period.
[42] Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).
[43] Alphabet Inc., 2023. 2023 Q1 10-Q, page 29 [accessed 26 September 2023]; Alphabet Inc., 2023. 2023 Q2 10-Q, page 32 [accessed 26 September 2023].

Figure A2.4: Annual global revenue for major cloud providers from their latest reported financial year (£bn)

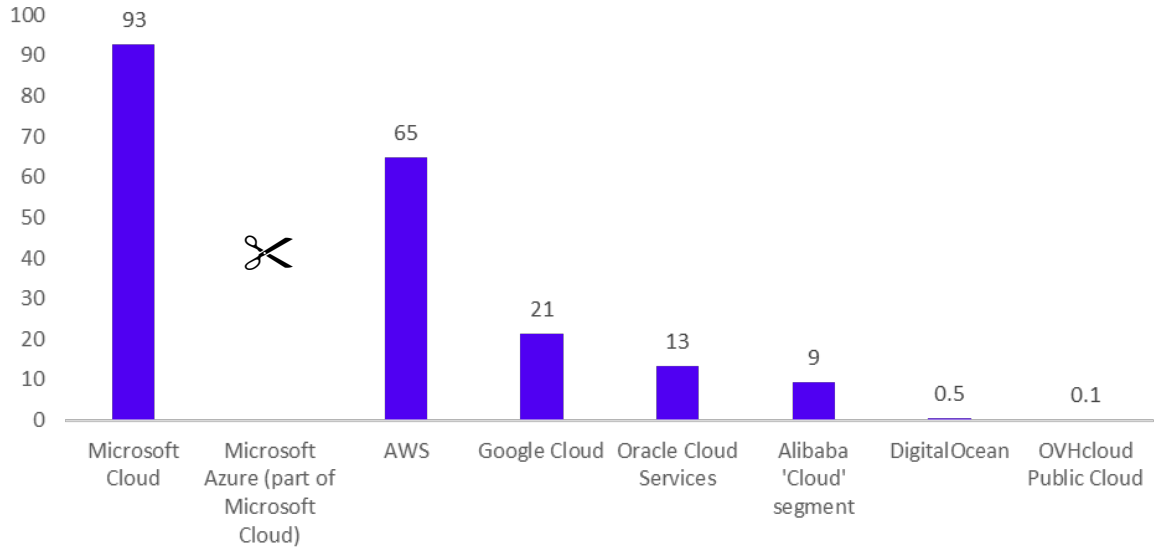

Source: Ofcom analysis of cloud providers' published financial statements and information provided by Microsoft in response to our information requests.[44] Figures (other than Azure) come from the following financial statements: Microsoft – year to June 2023, AWS – year to December 2022, Google Cloud – year to December 2022, Oracle Cloud Services – year to May 2023 (data on 'cloud services' taken from quarterly earnings releases), Alibaba 'Cloud' segment – year to March 2023, DigitalOcean – year to December 2022, OVHcloud Public Cloud – year to August 2022.

A2.18 IBM has, to date, reported revenues for 'Hybrid Cloud' in its annual report. Hybrid Cloud includes IBM's public cloud offerings, but also private cloud and hybrid cloud consulting and software solutions.[45] In 2022 IBM reported Hybrid Cloud revenue of $22.4bn, up 11% from $20.2bn in 2021. In response to our information request, IBM told us that in 2022 its global public cloud revenue was $[✂], which was lower than the hyperscalers.[46]

A2.19 This data demonstrates that global revenues for Microsoft Cloud and AWS are much higher than smaller cloud providers. As Microsoft Cloud includes revenue for Azure as well as other Microsoft cloud activities, Azure revenues will be lower than this. As indicated in the chart, in Microsoft's 2023 financial year, Azure represented around [✂] of Microsoft Cloud revenues. Google Cloud's most recent reported revenues were closer to those of smaller cloud providers like IBM and Oracle than AWS.

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

[45] Page 6 of IBM's 2022 Annual Report says the following about Hybrid Cloud: "*To provide useful decision-making information for management and shareholders, we define and measure hybrid cloud revenue as end-to-end cloud capabilities within hybrid cloud environments, which includes technology (software and hardware), services and solutions to enable clients to implement cloud solutions across public, private and multi-clouds. This spans across IBM's Consulting, Software and Infrastructure segments. Examples include (but are not limited to) Red Hat Enterprise Linux (RHEL), Red Hat OpenShift, Cloud Paks, as-a-service offerings, service engagements related to cloud deployment of technology and applications, and infrastructure used in cloud deployments*".

[46] IBM response dated 4 April 2023 to the s.174 notice dated 25 October 2022, Part B question 4 (Annex 2 submission).

A2.20 Figure A2.5 below shows quarterly year-on-year global cloud revenue growth for the hyperscalers, IBM Hybrid Cloud (which, as noted above, is a lot broader than IBM's public cloud offerings) and Oracle Cloud Services between 2018 and 2023.

Figure A2.5: Quarterly year-on-year global revenue growth for major cloud providers, 2018-23

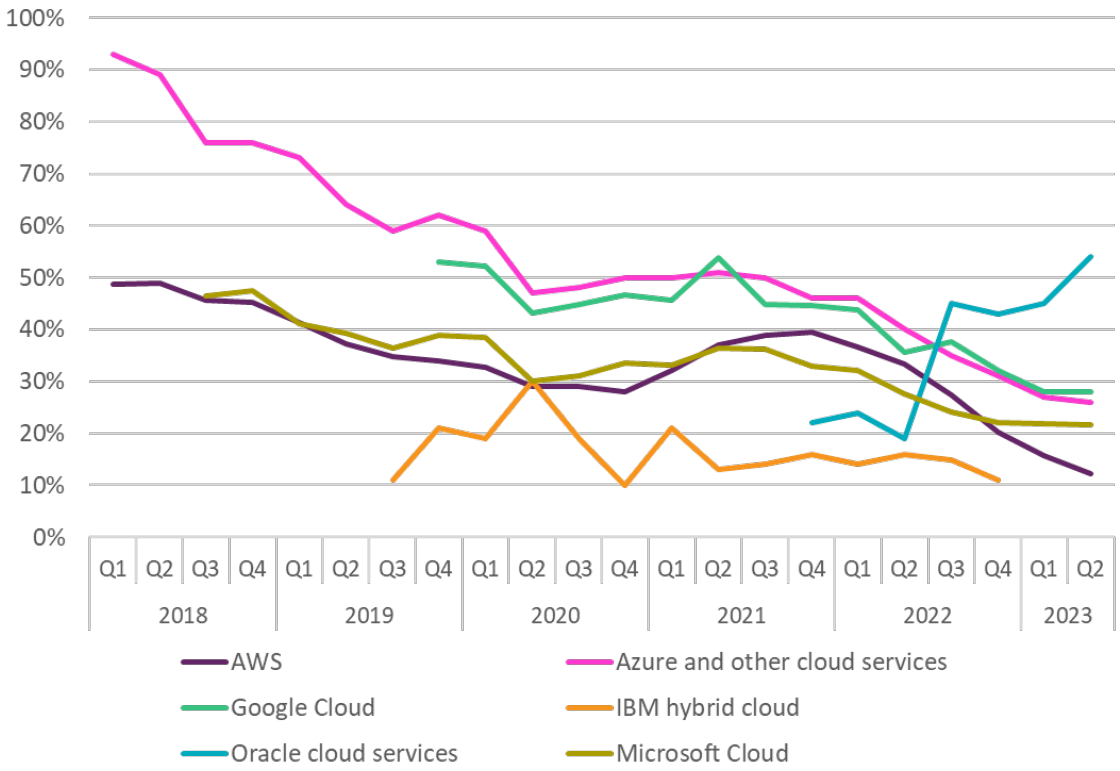

Source: Ofcom analysis of cloud providers' published financial statements.

Chart notes: Microsoft: Prior to the quarter ended September 2021, Microsoft reported revenue growth for Azure. Since then, it has reported revenue growth for 'Azure and other cloud services'. IBM: IBM reported 'total cloud revenues' from the quarter ended September 2019, then 'hybrid cloud revenues' from the quarter ended December 2021. IBM appears to have stopped reporting 'hybrid cloud revenues' from the quarter ended March 2023 onwards. Oracle: Oracle has reported cloud services revenue since the quarter ended August 2021. Oracle's financial year ends in May, but for the purposes of this chart we have used the closest available financial quarter corresponding to each calendar quarter.

A2.21 This data shows that the hyperscalers, IBM and Oracle have all recorded consistent double digit quarterly year-on-year revenue growth over this period. Comparing across providers, hyperscalers' global cloud revenue grew quicker than IBM and Oracle's cloud services for most of the last five years. However, this growth has been steadily declining and in recent quarters Oracle's cloud revenue increased at the fastest rate, albeit from a lower base.

A2.22 The chart also indicates that, over most of this period, Azure's revenue grew faster than the revenue of Microsoft Cloud and other cloud providers.

A2.23 We also obtained UK revenue estimates from cloud providers through our statutory information requests, which cover the period 2019 to 2022. This data suggests that during this period UK revenue growth was broadly consistent with global revenue growth.

Drivers of cloud revenue growth

A2.24 Growth in cloud revenues could be caused by:

a) increases in the unit prices customers are charged for consuming cloud services;
b) growth in the number of customers purchasing cloud services; and/or
c) growth in the usage of chargeable services consumed by existing customers.

A2.25 Based on responses to our statutory information requests, revenue growth appears to be driven (at least in part) by growth in the number of customers purchasing cloud services.[47] It is challenging to precisely quantify the extent of growth in the usage of chargeable services by existing customers, but based on the overall evidence available, we consider it is likely that existing customers' usage has grown and therefore contributed to revenue growth.[48]

A2.26 By contrast, data submitted by AWS and Microsoft suggests that average list prices for their core cloud infrastructure services have either remained stable or decreased in recent years, and that average net prices have decreased over the same timeframe.[49] This suggests that increases in unit prices have contributed less to growth in cloud revenues.

Trends in cloud profit margins

A2.27 In this sub-section we compare the operating profit (EBIT[50]) for hyperscaler cloud businesses to other cloud providers. AWS and Google Cloud publicly report EBIT. However, Microsoft does not publicly report EBIT for its Azure or Microsoft Cloud businesses. We have estimated EBIT for Microsoft Cloud based on publicly reported data, and the EBIT for Azure based on data provided by Microsoft in response to our statutory information requests. We describe below how we have done this, which is the same approach we used for estimating Microsoft Cloud EBIT and Azure EBIT in the interim report, although we have updated our analysis to take account of Microsoft's 2023 results. We did not receive any comments from stakeholders regarding our interim report approach.

Estimating Microsoft Cloud EBIT

A2.28 Since 2016, Microsoft has reported the gross margin percentage for Microsoft Cloud, and this margin includes an allocation of infrastructure costs such as data centres.[51] Microsoft Cloud gross margins have been around 70% in Microsoft's last four financial years (2020 to

[47] Our evidence on this comes from the UK. For example, all of the hyperscalers have grown their UK public cloud customer base in recent years – whether in terms of number of customers or number of customer accounts (where some customers could have multiple accounts). We note that AWS provided data on the number of UK AWS accounts rather than number of customers. Hyperscalers' information was provided via: AWS supplemental response dated 22 December 2022 to the s.174 notice dated 24 October 2022, Part B question 15; AWS response dated 31 March 2023 to the s.174 notice dated 24 October 2022, Part B question 15; Google response dated 16 December 2022 to the s.174 notice dated 26 October 2022, Part B question 15; Google response dated 31 March 2023 to the s.174 notice dated 26 October 2022, Part B question 15; Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, Part B question 15; Microsoft response dated 25 April 2023 to the s.174 notice dated 21 October 2022, Part B question 15.

[48] For example, as noted in Section 3, our customer research found that 79% of respondents expect to increase their spend on cloud services (relative to current levels) in the next 18 months. Source: Context Consulting research report, slide 42.

[49] Pricing trends for AWS and Microsoft Azure are further discussed in Section 4.

[50] EBIT stands for Earnings Before Interest and Tax.

[51] Microsoft's description of what is included in 'cost of revenue' (which is subtracted from revenue to estimate gross margin) says this includes data centre costs. See for example page 66 of its 2023 10-K.

2023). As Microsoft also reports revenue for Microsoft Cloud, we estimated the gross margin (in absolute terms) by multiplying revenue by the gross margin percentage.[52]

A2.29 To estimate EBIT for the financial years ending 2016 to 2023, we needed to attribute a proportion of Microsoft's remaining operating costs to Microsoft Cloud. Of these remaining operating costs, Research and Development and Sales and Marketing are the largest, representing around 87% of Microsoft's operating costs in its 2023 financial year, with General and Administrative costs representing around 13%.[53]

A2.30 Microsoft's 10-K[54] indicates that Microsoft generally allocates Sales and Marketing costs based on relative gross margin, but it does not specifically state the allocation basis used for other types of operating costs.[55]

A2.31 We considered attributing these costs based on revenues, cost of revenue and gross profits. The trend in EBIT derived using each of these approaches is very similar, with the EBIT derived from an allocation based on revenue generally being between the EBIT derived from the other approaches. We think it is reasonable to assume that Sales and Marketing and Research and Development costs will tend to scale with revenues.[56] As these represent the majority of remaining operating costs, we used revenue as the basis of allocation.[57]

Estimating Azure EBIT

A2.32 Microsoft provided us with an estimate of EBIT for Azure for its financial years ending 2018 to 2023 in response to our statutory information requests.[58] These estimates included an allocation of infrastructure costs to Azure but did not include an allocation of other operating costs.

A2.33 Microsoft also provided us with an estimate of EBIT (including allocations of all operating costs) for its 'Cloud and Enterprise' business, which is a subset of its Intelligent Cloud operating segment. Azure is a part of Cloud and Enterprise.[59]

---

[52] Our estimates of Microsoft Cloud revenue take account of Microsoft's recast of its revenue for the financial years ended 2017 and 2018 reflecting the addition of the commercial portion of LinkedIn to reported Microsoft Cloud revenue. This recast is reported in pages 32 and 91 of Microsoft's 2019 10-K, and it appears to have affected Microsoft Cloud revenue but not its reported gross margin percentage. We have not seen a recast estimate of Microsoft Cloud revenue for the financial year ended 2016, so our FY16 revenue estimate does not include revenue associated with the commercial portion of LinkedIn. We have taken account of this recast (where possible) in our estimation of both annual and quarterly Microsoft Cloud operating profit.

[53] Calculated based on figures presented on page 58 of Microsoft's 2023 10-K.

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

[55] Page 94 of Microsoft's 2023 10-K. This page also says some corporate costs are generally allocated to segments using gross margins or headcount.

[56] For example, to sustain higher revenues the company may need to invest in larger sales and marketing teams and continuous product development.

[57] We recognise there could be elements of these cost categories which may be more or less relevant to Microsoft's cloud activities. For example, page 17 of Microsoft's 2023 10-K says that there are various engineering groups within Research and Development, some of which appear more relevant to cloud than others. However, we do not have data available to take account of these potential effects.

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

[59] Microsoft response dated 20 January 2023 to our follow-up email dated 20 December 2022 concerning the s.174 notice dated 21 October 2022, Part B question 22; Microsoft response dated 4 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).

A2.34 We asked Microsoft to explain how costs are allocated to its cloud businesses, but it did not explain its allocation approach in detail.[60] We therefore needed to make some assumptions to estimate an EBIT for Azure. We estimated EBIT for Azure using the same approach as for Microsoft Cloud, i.e. we allocated remaining operating costs to Azure on the basis of relative revenue. We note that, based on the data provided by Microsoft, the share of remaining Microsoft operating costs for Cloud and Enterprise is very similar to its share of total Microsoft revenue.[61] This suggests that allocating remaining operating costs to Azure using revenue could give a reasonable estimate of its EBIT, given that Azure is part of Cloud and Enterprise.

A2.35 As for Microsoft Cloud, we also considered allocating remaining operating costs to Azure based on relative share of gross profits and cost of revenue. The trend in EBIT derived using each of these approaches is similar, with the EBIT derived from an allocation based on revenue generally being between the EBIT derived from the other approaches. Therefore, the choice of allocation approach is unlikely to affect our findings.

Comparison of EBIT with other cloud providers

A2.36 Figure A2.6 below shows the latest annual EBIT reported for AWS and Google Cloud alongside our estimates of EBIT for Microsoft Cloud and Azure (which has been redacted in the published version of this report) and compares this to the EBIT reported for Alibaba's reported cloud segment and DigitalOcean. 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.[62]

A2.37 This chart shows that in absolute terms, the most recent annual EBIT for AWS (£18bn) and our estimate of EBIT for Microsoft Cloud (£42bn) are significantly higher than the EBIT of other cloud providers, including Google Cloud. In its most recent financial year, the EBIT for Google Cloud was -£2bn and on an annual basis it has been consistently loss making to date, although it reported its first quarterly operating profit of £157m in the quarter ending March 2023.[63] We estimate that Azure EBIT in Microsoft's 2023 financial year was [✂]. This is [✂].

[60] 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.
[61] Microsoft response dated 20 January 2023 to our follow-up email dated 20 December 2022 concerning the s.174 notice dated 21 October 2022, Part B question 22; Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).
[62] IBM response dated 23 December 2022 to our follow-up email dated 9 December 2022 concerning the s.174 notice dated 25 October 2022, Part B question 9; Oracle response dated 13 January 2023 to questions 7 and 14 of our follow-up email dated 22 December 2022 concerning the s.174 notice dated 31 October 2022, Part B questions 9 and 20; Oracle response dated 2 June 2023 to our follow-up email dated 10 May 2023 concerning the s.174 notice dated 31 October 2022, Part B question 9.
[63] Alphabet Inc., 2023. 2023 Q1 10-Q, page 29 [accessed 26 September 2023]. In the quarter ended June 2023 Google Cloud reported an EBIT of £316m (Alphabet Inc., 2023. 2023 Q2 earnings release, page 2 [accessed 4 September 2023]. We have used the ONS Average Sterling exchange rate: US Dollar series to currency convert Google Cloud's reported quarterly operating profits from US dollars to pound sterling.

Figure A2.6: Annual global EBIT for the latest financial year (£bn)

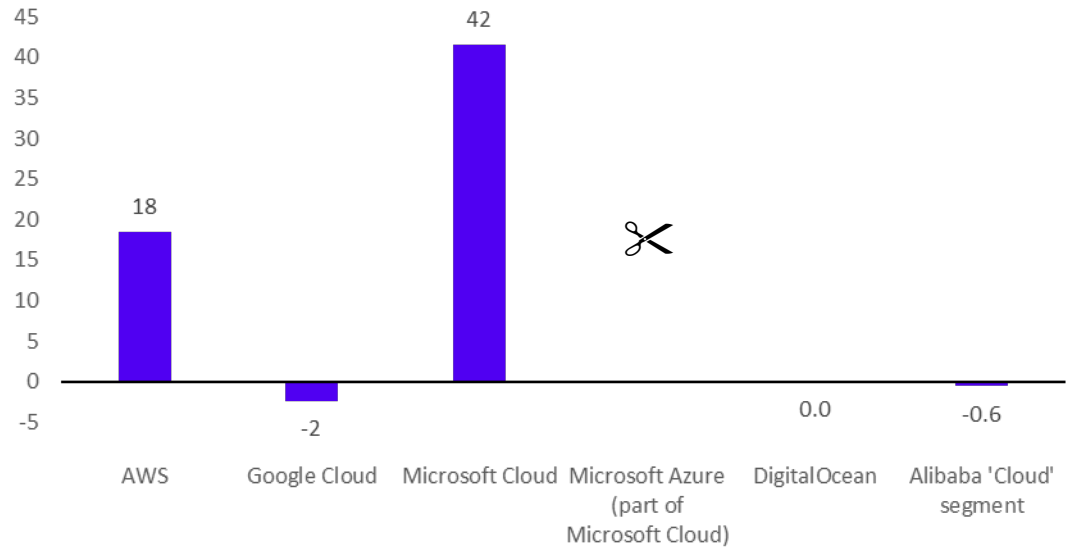

Source: Ofcom analysis of cloud providers' published financial statements, information provided by Microsoft in response to our information requests[64] and Ofcom assumptions. Figures (other than Azure) come from the following financial statements: Microsoft – year to June 2023, AWS – year to December 2022, Google Cloud – year to December 2022, Alibaba 'Cloud' segment – year to March 2023, DigitalOcean – year to December 2022.

A2.38 Figure A2.7 below captures quarterly EBIT margins between 2018 and 2023 for AWS, Google Cloud, Alibaba's 'Cloud' segment and DigitalOcean.[65] It also includes our EBIT margin estimate for Microsoft Cloud and Azure following the methodology described above (though we only have information to estimate the quarterly EBIT for Microsoft Cloud since Q3 2020[66] and an annual EBIT margin for Azure). Our estimates for Azure have been redacted in the published version of this report. This chart includes two more quarters of data compared to the interim report.

[64] Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).
[65] Operating profit (EBIT) margin is calculated as operating profit divided by revenue.
[66] We can estimate annual EBIT for Microsoft Cloud prior to Q3 2020. We have therefore captured annual EBIT estimates for Microsoft's financial years ending 2018-2020 and we have plotted a linear trendline between each of these datapoints to ensure visibility.

Figure A2.7: Quarterly EBIT margins for major cloud providers, 2018-23

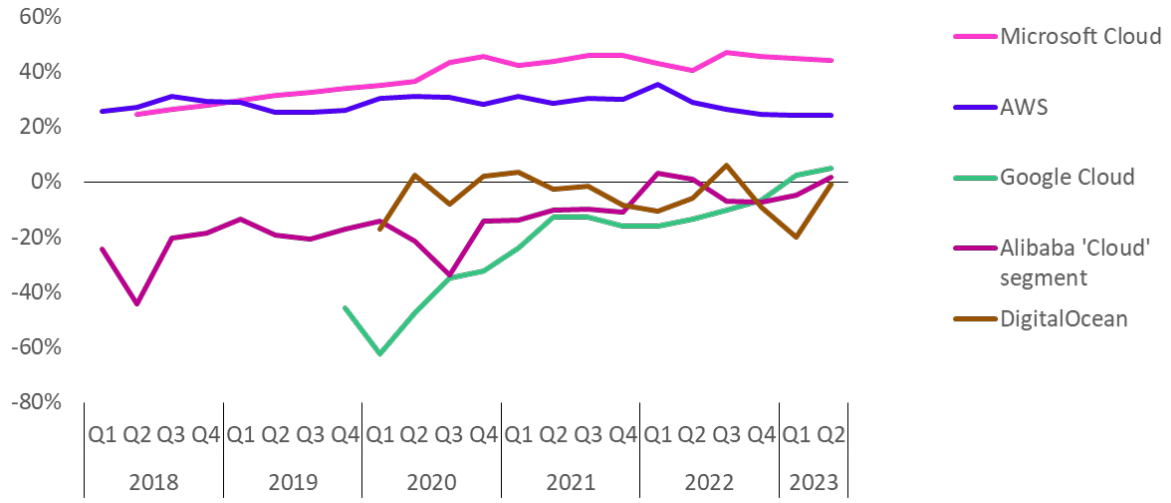

Source: Ofcom analysis of cloud providers' financial data reported by S&P Capital IQ, information provided by Microsoft in response to our information requests[67] and Ofcom assumptions. We have used the latest filings (incorporating restatements).

A2.39 OVHcloud's 'Public Cloud' segment had an EBITDA margin[68] of 35% to 40% in its 2021 and 2022 financial years respectively.[69] OVHcloud does not report EBIT margins for its Public Cloud segment, but we note EBIT margins for its overall business were close to zero in these years.[70]

A2.40 For the hyperscalers, we estimate that, over this period:

a) AWS had stable EBIT margins of around 20% to 30%.
b) 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 [✂] in the year to June 2023 compared to 25% for AWS.
c) Google Cloud made losses until the end of 2022, but it has since reported operating profits for the first two quarters of 2023, and its EBIT margins are trending upwards, though remain low compared to AWS and Microsoft Cloud. Google Cloud's recent EBIT margins appear to have been positively impacted by accounting policy changes Alphabet made in 2023 (namely an increase in the estimated useful lives of servers and certain network equipment, and changes in segment cost allocations).[71] We observe that when AWS was generating similar revenues to Google Cloud, its reported EBIT margins were

[67] Microsoft response dated 20 January 2023 to our follow-up email dated 20 December 2022 concerning the s.174 notice dated 21 October 2022, Part B question 22; Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).
[68] EBITDA stands for Earnings Before Interest, Tax, Depreciation and Amortisation.
[69] Calculated from OVHcloud 2022 Financial Statements, page 27.
[70] OVHcloud's overall business had an EBITDA of around 35% in these years, but after depreciation and amortisation, its EBIT margin was around 1% to 2%. OVHcloud's Public Cloud segment is around 16% of group revenues.
[71] These accounting policy changes took effect from the first quarter of 2023. For more information see: Alphabet Inc., 2023. An update on certain reporting and disclosure topics in our Q1 2023 earnings [accessed 26 September 2023]; Alphabet Inc., 2023. Alphabet Announces First Quarter 2023 Results [accessed 26 September 2023].

between 25% and 29%. This implies Google Cloud's costs are higher than AWS's when it was in a similar revenue position.[72]

A2.41 Among other cloud providers, quarterly EBIT margins for Alibaba's 'Cloud' segment and DigitalOcean were occasionally positive but mostly negative in recent years. Lower profits observed for smaller providers could be consistent with a lack of current scale, especially where investment is required ahead of growing revenues (e.g. in data centre capacity, sales teams and product offerings).

A2.42 We also received additional insights from responses to our statutory information requests.

a) Information provided by IBM indicates that its global public cloud business [✂]. IBM said that the need to have a large global infrastructure footprint combined with sufficient scale are the two most important factors influencing utilisation in datacentres and profitability.[73]
b) Oracle said it could not provide detailed or accurate profit estimates on cloud beyond its public reporting for the Cloud and License segment, though it did say that its SaaS services were the most profitable, followed by PaaS and then IaaS.[74]

A2.43 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 reducing losses) among smaller cloud providers, these are a lot lower than those for AWS and Microsoft Cloud. Azure operating profits [✂].

Assessment of ROCE relative to WACC

Why we use ROCE as a measure of profitability

A2.44 Evaluation of ROCE relative to the cost of capital is a standard approach to assessing firms' profitability, which is highlighted in the CMA's Guidelines for market investigations[75] and has been frequently used in past CMA market studies and Ofcom profitability assessments.[76]

A2.45 ROCE can be compared against the WACC to test whether profits are high. The CMA has previously noted that, while companies need to earn positive margins to be sustainable, margins on their own do not indicate whether profit is higher than might be expected in a market that is working well - some sectors with high asset investment and low operating

[72] In Q2 2023 Google Cloud revenue was $8.03bn, as reported in page 2 of Alphabet's Q2 2023 results press release. Our analysis of quarterly AWS revenue shows that AWS revenue was closest to this between Q1 2019 and Q2 2019 when AWS generated revenue of $7.70bn (Q1 2019) and $8.38bn (Q2 2019), as reported in page 18 of Amazon's Q1 2019 10-Q report and page 18 of Amazon's Q2 2019 10-Q report. Combined with the reported AWS operating income of $2.22bn (Q1 2019) and $2.12bn (Q2 2019), this tells us that AWS EBIT margin was 28.9% in Q1 2019 and 25.3% in Q2 2019.
[73] 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.
[74] Oracle response dated 6 January 2023 to our follow-up email dated 22 December 2022 concerning the s.174 notice dated 31 October 2022, Part B question 9.
[75] Competition Commission, 2013. Guidelines for market investigations: Their role, procedures, assessment and remedies, Annex A paragraph 9 [accessed 10 September 2023].
[76] See for example: (1) CMA, 2022. Appendix D: financial analysis of Apple's and Google's mobile ecosystems [accessed 10 September 2023]; (2) Ofcom, 2022. Annex: Ofcom's future approach to mobile markets [accessed 10 September 2023].

costs will have high margins, but that would not necessarily equate to high economic profitability. ROCE can be compared against the profit a company would require to recover the cost of investments made in the past.[77]

A2.46 A finding that ROCE is higher than the WACC in any particular period is not necessarily indicative of a competition problem. A firm that innovates and gains a competitive advantage may earn higher ROCE for the period that it is able to sustain that competitive advantage. In a market characterised by effective competition, we would expect returns to converge towards WACC over time, as competitors would enter the market to benefit from the high returns on capital and compete away the difference between ROCE and WACC. This also applies to markets with significant levels of innovation, as the high returns on capital provide an incentive for competitors to enter the market with similar or competing technologies.

A2.47 The CMA's guidance therefore says a finding that 'profitability of firms representing a substantial part of the market has exceeded the cost of capital over a sustained period could be an indication of limitations in the competitive process'.[78] As this guidance implies, profitability analysis is only an indicator of limitations in competition and cannot alone provide conclusive evidence around the level of competition in a market.

A2.48 In response to our interim report, Microsoft suggested that ROCE and WACC cannot be compared, as WACC has been developed as a discount rate for after-tax cash flows, whereas ROCE uses EBIT as a measure of profits, which is pre-tax and does not include some cash costs.[79]

A2.49 We disagree with Microsoft that ROCE and WACC cannot be compared. Evaluating ROCE against WACC is a standard approach for assessing profitability, which has been used frequently in past CMA and Ofcom work. We agree that, where detailed cashflow data is available, it may be appropriate to use a cashflow-based analysis such as internal rate of return (IRR) to assess profitability.[80] However, as noted in the CMA's Guidelines, the choice of profitability measure partly depends on data availability.[81] As we do not have detailed cashflow data for AWS, Azure or Microsoft Cloud, we assessed profitability using ROCE.

A2.50 We agree that a consistent approach to tax needs to be taken. Our analysis takes account of tax considerations by comparing ROCE calculated using EBIT against pre-tax WACC estimates which are adjusted for tax using evidence of effective tax rates paid by Amazon and Microsoft.[82]

---

[77] CMA, 2020. Appendix D: profitability of Google and Facebook, paragraphs 8 and 9 [accessed 10 September 2023].
[78] Competition Commission, 2013. Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraph 118 [accessed 10 September 2023].
[79] Microsoft response to the interim report, page 39, paragraph 152.
[80] For example, as used by Ofcom in Pay TV market investigation in 2011 and the CMA in its recent Mobile Radio Network Services market investigation. The CMA Guidelines also say it may be appropriate to use IRR where reliable data is available. Source: Competition Commission, 2013. Guidelines for market investigations: Their role, procedures, assessment and remedies, Annex A, paragraph 10 [accessed 10 September 2023].
[81] Competition Commission, 2013. Guidelines for market investigations: Their role, procedures, assessment and remedies, Annex A paragraph 11 [accessed 10 September 2023].
[82] This is consistent with the CMA's approach in recent market studies – e.g. its 2020 market study into online platforms and digital advertising, where it compared ROCE calculated using EBIT to a pre-tax WACC. See: CMA, 2020. Appendix D: profitability of Google and Facebook [accessed 26 September 2023].

A2.51 Finally, we recognise that EBIT is influenced by accounting choices, particularly in relation to depreciation of assets. For example, in recent years hyperscalers have increased the assumed useful life of servers, which has generally had the effect of increasing EBIT. The impact of accounting choices on particular periods is one reason why it is important to consider ROCE over a relatively long period, as we have endeavoured to do, to test for persistency of returns above WACC.

Our approach to estimating ROCE

A2.52 ROCE is calculated by dividing EBIT by the value of capital that is employed in the relevant business. EBIT and capital employed are derived from accounting data, but adjustments may be required to estimate an economically meaningful measure of profitability. The principle of these adjustments is to ensure that profits and capital employed arising from the operation of the relevant business are included in the analysis. This means that financing costs (such as cash and other sources of finance) and tax costs are generally excluded.[83]

A2.53 For EBIT, we have relied on our estimates for AWS and Microsoft set out above. We considered whether any adjustments were required to remove any one-off exceptional expenses which may distort our estimates of EBIT, such as the costs of legal settlements, but we did not identify any relevant costs. We therefore used our estimated EBITs.

A2.54 Capital employed should represent a reasonable estimate of the replacement cost of assets and liabilities required to operate the relevant business, which may be different from accounting values recorded in financial statements.[84] There is a degree of judgement about which assets and liabilities should be included in calculating capital employed. We have tested different approaches in our ROCE analysis, but our general approach reflects the following:

a) **Non-current assets**: We consider it appropriate to include property and equipment assets (and associated leases) used by the relevant business, as these assets are often directly linked to provision of the relevant services. It may be appropriate to include other non-current assets where these are necessary to provide the relevant services (for example intangible assets), however, we exclude goodwill from capital employed. Goodwill is not a separately identifiable asset but a balancing figure between the purchase price and the fair value of acquired assets. Including goodwill risks capitalising the future value of any excess profits that a business can generate, which may be reflected in the purchase price and hence the purchased goodwill.[85]
b) **Current assets and liabilities**: We exclude cash and marketable securities from current assets, as these are more relevant to financing policy choices than the operation of the

[83] This is consistent with the CMA's approach in previous market studies, such as its 2020 market study into online platforms and digital advertising. See: CMA, 2020. Appendix D: profitability of Google and Facebook, paragraph 19 [accessed 10 September 2023].
[84] The CMA guidelines for market investigations says it considers replacement cost the economically meaningful measure of capital employed in most cases. See Competition Commission, 2013. Guidelines for market investigations: Their role, procedures, assessment and remedies, Annex A paragraph 14 [accessed 26 September 2023].
[85] Our treatment of goodwill is consistent with Ofcom's approach to other ROCE analyses, such as our 2022 discussion paper on Ofcom's future approach to mobile markets. See: Ofcom, 2022. Annex: Ofcom's future approach to mobile markets, paragraph A6.10 [accessed 10 September 2023]. Our approach is also informed by previous CMA analyses of ROCE, such as its 2016 energy market investigation. See: CMA, 2016. Appendix 9.10: Analysis of retail supply profitability – ROCE, paragraphs 60-61 [accessed 10 September 2023].

relevant business.[86] It may be appropriate to include other current assets and liabilities, such as accounts receivable and accounts payable, in capital employed where these are necessary to deliver the relevant services.

Scope of our cloud ROCE analysis

A2.55 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. For context and comparison with Azure, we also present an estimate of ROCE for Microsoft Cloud. We have not included Google Cloud in our ROCE analysis as, to date, it has reported annual operating losses (so its ROCE calculated using annual data in recent years will be negative and below WACC).[87]

A2.56 Our analysis draws on global data from AWS and Microsoft, some of which is publicly reported and some of which has been provided to us in response to our statutory information requests. To help us assess capital employed, we requested global balance sheet data for AWS, Azure and Microsoft Cloud from AWS and Microsoft. However, AWS was unable to provide any more data on AWS assets than that presented in Amazon's annual reports, and Microsoft could not provide any balance sheet data for Azure or Microsoft Cloud. We have therefore had to make assumptions when estimating capital employed, as set out below.

A2.57 For each of Amazon and Microsoft, we explain below how we have estimated ROCE. Our approach is consistent with that presented in our interim report, although we have updated our analysis to take account of Microsoft's 2023 financial results. We did not receive any detailed comments from stakeholders relating to how we estimated ROCE or the size of our ROCE estimates.

A2.58 The rest of this sub-section is structured as follows:

a) We first estimate the WACC applicable to cloud activities.
b) We then explain our assessment of ROCE applicable to AWS, Microsoft Azure and Microsoft Cloud, and compare ROCE to WACC.

WACC applicable to cloud activities

A2.59 Our approach to estimating WACC reflects the circumstances of this market study. It should not be interpreted as an indication of how we might evaluate the WACC for a different purpose, such as setting prices applicable to a regulated business.

A2.60 We asked AWS and Microsoft if they estimated WACC for AWS, Azure and Microsoft Cloud, but both said they did not produce WACC estimates at this level.[88] In our interim report, we therefore estimated a WACC to capture the systematic risk associated with cloud activities. However, since we published our interim report Microsoft has told us that [✂] prepares

[86] Excluding cash and marketable securities is consistent with the CMA's approach in previous market studies, such as its 2020 market study into online platforms and digital advertising and its 2022 market study into mobile ecosystems.
[87] While Google Cloud reported operating profits for the first two quarters of 2023, its operating margin was low compared to AWS and Microsoft as shown in Figure A2.7.
[88] AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, Part B question 29; Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, Part B question 29; Microsoft response dated 31 May 2023 to the s.174 notice dated 23 May 2023, question 3.

WACC estimates for its Cloud and Enterprise business (which includes Azure) for the purposes of annual goodwill impairment testing,[89] which we consider below as a sense-check of our WACC estimate.

A2.61 We estimated that the pre-tax nominal WACC applicable to AWS and Azure for the period we are considering was likely to fall between 9.0% and 13.0%. We think this range remains reasonable for the following reasons:

a) The CMA estimated a pre-tax nominal WACC of 9% for Google and Facebook as part of its 2020 market study into online platforms and digital advertising.[90]
b) When estimating the WACC for Google and Facebook, the CMA used a beta range of 1.0 to 1.15.[91] This range appears to capture recent asset beta estimates for the hyperscalers[92] and the lower end of the range is broadly consistent with the asset betas for cloud computing exchange traded funds.[93] However, we observe that asset betas for the hyperscalers were somewhat higher before the Covid lockdowns in 2020. For example, Amazon and Microsoft asset betas were mostly between 1.15 and 1.3 in 2019. As we are looking at returns over several years, and as the WACC should reflect investor expectations at the time of investment, we have used an asset beta range of 1.0 to 1.30 to reflect the possibility that cloud betas could have been higher in the past.[94]
c) The CMA applied an effective tax rate of 9% in its calculation for Google and Facebook, which was specific to these businesses. In recent years the effective tax rates paid by Amazon and Microsoft have averaged around 15%, with some evidence that they were higher before then. We have therefore used an effective tax rate range of 15% to 20%.
d) Combining our beta range and effective tax rate range with the market parameters used by the CMA in its WACC estimate for Google and Facebook, we obtain a WACC range of 8.9% to 13.2%.[95] As the high end of our beta and effective tax range reflects older data, the high end of our WACC range is arguably more relevant to investments made some years ago, with the low end of the range relevant to more recent investments.

A2.62 As a sense check, we compared our WACC range to Microsoft's internal WACC estimates for its Cloud and Enterprise business, of which Azure is a part. Between 2018 and 2022 these

---

[89] Microsoft response dated 31 May 2023 to the s.174 notice dated 23 May 2023, question 3.
[90] CMA, 2020. Appendix D: profitability of Google and Facebook, pages D39 to D41 [accessed 10 September 2023].
[91] As the CMA's WACC estimate was based on the pre-tax cost of equity, assuming no debt, we assume these betas represent asset betas. The asset beta is calculated by un-levering the equity beta for the effect of gearing. As companies like Amazon and Microsoft have relatively low gearing the equity and asset betas are similar.
[92] By reference to 5-year daily asset betas as at June 2023, measured against the S&P 500 index. Asset betas are estimated using the following formula assuming a debt beta of 0.1: $\beta_a = \beta_e * (1 - gearing) + gearing * \beta_d$. The hyperscalers' betas will reflect the systematic risk of their overall business, not just cloud, though cloud represents an increasing proportion of overall activity. For example, in recent years AWS has contributed a significant proportion of Amazon's operating profits (close to 90% on average between 2018 and 2022).
[93] Such as First Trust Cloud Computing ETF (ticker SKYY) and Global X Cloud Computing (CLOU). These funds aim to track the performance of companies involved in the cloud computing industry. We estimate equity betas of these funds of around 1.10 – 1.15 in June 2023. Many of the companies tracked by these funds have relatively low gearing, so it is likely that the corresponding asset beta of these funds will be around 1.0.
[94] Though we recognise that cloud was a smaller part of these companies in the past, and as such these higher betas may not reflect the systematic risk of their cloud activities.
[95] As per the CMA's estimate for Google and Facebook, for the purposes of this estimate, WACC is equal to the pre-tax cost of equity. For much of the period we are considering, Amazon and Microsoft had relatively low gearing, so we have not included a cost of debt.

estimates are broadly consistent with our WACC range.[96] We also compared our WACC range against the 'Rest of BT' (RoBT) WACC we estimated in previous telecoms market reviews. Since 2016, the RoBT WACC has reflected the systematic risk associated with BT's ICT activities,[97] and it could be a relevant benchmark to the extent the systematic risk associated with cloud activities is similar to that for global ICT services. In 2016, our RoBT pre-tax nominal WACC estimate was around 13%[98] and in subsequent assessments it generally reduced, falling to around 10% in our most recent decision in 2021.[99] On this basis, our estimated range of 9% to 13% for cloud does not appear unreasonable.

A2.63 In response to our interim report, AWS said our beta estimate is flawed as there is no comparable firm of its size in the US that operates only as a cloud provider or in the provision of cloud infrastructure services.[100] AWS did not suggest an alternative beta estimate. Microsoft said that the estimation of ROCE and WACC should take account of 'asset premia due to capacity constraints'.[101] Microsoft did not explain what this meant or how we should take account of it, but we have assumed Microsoft was suggesting the WACC should include a premium to reflect the possibility that capacity constraints could limit upside returns.

A2.64 While we agree there are no 'pure-play' cloud providers against which to benchmark beta, we have used a wide beta range of 1.0 to 1.30 informed by beta estimates for the hyperscalers and cloud computing exchange traded funds. We think this gives a reasonable view of the likely range for beta. We note that the betas used in Microsoft's internal WACC estimates for its Cloud and Enterprise business (which includes Azure) over the period 2018 to 2022 were generally towards the lower end of our beta range, and the highest hyperscaler betas in the last few years did not exceed the top end of our range. To the extent the systematic risk of cloud services could be affected by factors such as capacity constraints, we consider our beta range is wide enough to capture these effects.

A2.65 As set out above, we have also sense checked our resulting WACC range against two benchmarks, both of which indicate our WACC range is reasonable.

A2.66 Microsoft also said that we have used a static value for WACC when assessing ROCE, and that this approach is not suitable because WACC is a dynamic measure.[102] We disagree we have used static WACC – our range takes account of how WACC parameters such as beta and effective tax rates have evolved over time, and incorporates parameter values seen in past years (for example, higher betas observed prior to the Covid lockdowns in 2020). We therefore consider our range reflects how WACC could have varied over the period for which we are assessing ROCE.

---

[96] [✂]. This is comparable to our 9% to 13% WACC range.
[97] BT's ICT operations are spread across its BT Global and other divisions. These include global security, cloud and networking services.
[98] Ofcom, 2016. Business Connectivity Market Review 2016, Annex 30, Table A30.2 [accessed 10 September 2023].
[99] Ofcom, 2021. Wholesale Fixed Telecoms Market Review 2021-26, Annex 21, Table A21.2 [accessed 10 September 2023].
[100] [✂].
[101] Microsoft response to the interim report, page 38, paragraph 150.
[102] Microsoft response to the interim report, pages 38-39, paragraph 151.

Our assessment of ROCE for AWS

A2.67 Amazon has separately reported financial information for AWS since 2013, including information on revenue, operating profit, net property and equipment assets, and total assets. The reported data allows us to estimate ROCE for AWS.

A2.68 To derive our baseline ROCE estimates for AWS, we used AWS's reported EBIT.

A2.69 To derive capital employed for AWS we considered the information Amazon publishes in its annual reports:

a) **AWS net property and equipment**: Amazon describes this as including buildings and land (which we assume includes owned data centres) as well as equipment such as servers and networking equipment.[103] This category also includes finance leases (which could include data centres owned under leasing arrangements).[104]
b) **AWS total assets**: Amazon says AWS total assets primarily consist of property and equipment assets and accounts receivable.[105] Amazon also says that the majority of technology infrastructure assets are allocated to AWS based on usage.[106] AWS total assets exclude corporate assets, such as cash and cash equivalents, marketable securities, other long-term investments, corporate facilities, goodwill and other acquired intangible assets, and tax assets.[107]

A2.70 As per our interim report, our baseline capital employed for AWS is based on net property and equipment assets. This represents investment in tangible assets required to operate AWS. It also includes finance leases, which are likely to include leased data centres relevant to AWS, which are required to provide relevant cloud services. While this measure does not include operating leases, Amazon's public reporting implies that these may be more relevant to its retail operations than AWS.[108] While this measure of capital employed does not include working capital, Amazon has had a net current liability balance for many years (excluding cash and equivalents consistent with our methodology above),[109] such that including these is likely to reduce capital employed and increase ROCE. On balance therefore, we think using net property and equipment will give a reasonable estimate of AWS capital employed.

A2.71 Figure A2.8 shows AWS EBIT and net property and equipment assets between 2013 and 2022. Both have increased significantly over this period. EBIT has generally grown faster than net property and equipment assets, so we would expect AWS ROCE based on this measure of capital employed to have increased over this period.

---

[103] Amazon, 2023. 2022 10-K, page 47 [accessed 10 September 2023].
[104] Amazon, 2023. 2022 10-K, page 47 [accessed 10 September 2023].
[105] Amazon, 2023. 2022 10-K, page 68 [accessed 10 September 2023].
[106] Amazon, 2023. 2022 10-K, page 68 [accessed 10 September 2023]. We understand technology infrastructure includes servers, networking equipment and data centres.
[107] Amazon, 2023. 2022 10-K, page 68 [accessed 10 September 2023].
[108] Page 68 of Amazon's 2022 10-K says AWS assets mostly consist of net property and equipment and accounts receivable, and that retail assets in its North America and International segments primarily consist of property and equipment, operating leases, inventory and accounts receivable - i.e. operating leases are noted as important for retail but not AWS.
[109] For example, excluding cash and cash equivalents, Amazon's current assets in 2022 were $77bn and its current liabilities were $155bn (a net current liability of $78bn).

Figure A2.8: AWS global EBIT and net property and equipment assets, 2013-2022, $bn

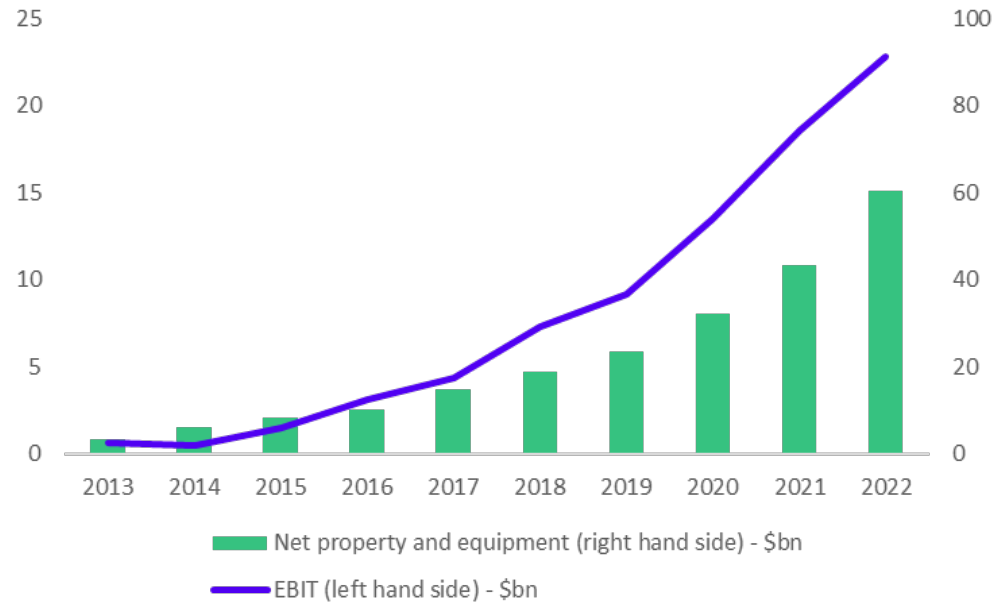

Source: Ofcom analysis based on Amazon 10-K reports.

A2.72 We recognise that using net property and equipment may not capture other assets relevant to AWS. We considered whether including these, along with relevant current liabilities, could affect our analysis using the following two sensitivities:

a) **Sensitivity 1: AWS total assets less an estimate of accounts payable**. We understand that AWS total assets (as reported in Amazon's annual report) include property and equipment (including finance leases), operating leases, accounts receivable and some other current and non-current assets (though we have no visibility of these other assets). Usually we would exclude cash, marketable securities and goodwill from capital employed, but AWS total assets already exclude these items. This measure of total assets will likely overstate capital employed (and understate ROCE) as it includes accounts receivable but not accounts payable (or any other relevant current liabilities). It may also include some assets which are not needed to provide the relevant cloud services. While we do not have the data to adjust the asset base to remove assets that may not be necessary to provide relevant cloud services, we have included an estimate of accounts payable. We estimated this for AWS by calculating accounts payable as a proportion of revenue for Amazon overall and applying this proportion to AWS revenue.

b) **Sensitivity 2: Include all technology infrastructure assets with AWS capital employed.** Amazon allocates technology infrastructure assets to AWS based on usage. AWS told us that these assets predominantly relate to data centres, networks and servers operated by AWS to deliver services to both AWS external customers, and internal Amazon non-AWS teams. The allocation of these assets to AWS reflects usage by external customers.[110] As a single data centre could be used by both external and internal customers, we have considered the effect of assuming all these technology assets are necessary to serve external customers, by allocating all of Amazon's technology infrastructure assets to AWS when calculating capital employed. We consider this is a conservative assumption as, in practice, the relevant technology infrastructure assets

[110] AWS response dated 13 January 2023 to question 8 of our follow-up email dated 16 December 2022 concerning the s.174 notice dated 24 October 2022, Part B, question 31.

required to serve external customers only may consist of smaller data centres or fewer servers. We are only able to make this adjustment for the period 2017 to 2022, based on information provided by AWS.[111]

A2.73 As noted above, AWS total assets exclude corporate assets. It is possible that some of these, such as corporate facilities and intangible assets, could be relevant to the estimation of AWS capital employed. Microsoft said that the estimation of ROCE should take account of intangible assets.[112] We consider Microsoft's comment below when estimating Azure ROCE, but we also consider whether the inclusion of intangible assets could affect our findings for AWS. We do not have visibility of the value of intangible assets for AWS, so we have not made a separate adjustment for this. However, based on the results of our wider analysis we consider it is unlikely that any adjustment would be substantial enough to affect our overall analysis. For example, we estimate from information provided by AWS that cloud-related acquisitions undertaken between 2017 and 2021 represented [✂].[113] This implies that intangible assets resulting from cloud-related acquisitions are likely to have a limited impact on our estimates of AWS EBIT and capital employed.

A2.74 We also recognise that the replacement costs of AWS's assets may be different from historical values. AWS could not provide us with a breakdown of its assets beyond those reported publicly but, as noted above, we know that net property and equipment assets include servers and networking equipment which are depreciated over five and six years respectively (four and five years prior to January 2022).[114] This suggests that, at least for these assets, there may not be a significant difference between the net book value and depreciated replacement cost, given the relatively short useful lives. Indeed, it is possible that the replacement cost for some assets, such as computing assets, is lower than historical cost (which could mean our estimate of capital employed is overstated for some assets). For other assets included in net property and equipment, such as datacentres, which have longer useful lives, there may be more of a difference between net book value and depreciated replacement cost. However, to the extent that our analysis has understated replacement costs, the margin of error would need to be large for our findings relating to AWS ROCE and WACC to be materially impacted. For example, we estimate that in 2022 our estimate of AWS capital employed would need to be almost three times as large for our baseline ROCE estimate to be comparable to the high end of our WACC range.

A2.75 It is also possible that some assets relevant to the provision of AWS services have already been fully depreciated. We asked AWS for the value of technology infrastructure assets which had been fully depreciated by 2013 which were used to provide AWS services or are still used to provide AWS services today. [✂].[115] However, given the trend in AWS net property and equipment assets shown in Figure A2.8, we think the value of any relevant

[111] AWS response dated 17 January 2023 to question 8 of our follow-up email dated 16 December 2022 concerning the s.174 notice dated 24 October 2022, Part B question 31. AWS response dated 26 May 2023 to question 6 of our follow-up email dated 5 May 2023 concerning the s.174 notice dated 24 October 2022, Part B question 31. We have not separately adjusted depreciation associated with technology infrastructure assets for the purposes of this sensitivity, though we do not think doing so would significantly affect the results.
[112] Microsoft response to the interim report, page 38, paragraph 150.
[113] AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, Part B question 27. We compared [✂].
[114] Amazon, 2023. 2022 10-K, page 41 [accessed 26 September 2023].
[115] AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, Part B question 31.

assets fully depreciated by 2013 could be relatively low compared to the asset base today, and so unlikely to materially affect our analysis.

A2.76 Figure A2.9 below shows our AWS ROCE estimates using our baseline approach and first sensitivity (based on total AWS assets less accounts payable). The data covers Amazon's financial years ending 2013 to 2022 (Amazon's financial year ends in December). This chart is the same as that presented in our interim report.

Figure A2.9: AWS ROCE estimates (baseline approach and first sensitivity), 2013-22

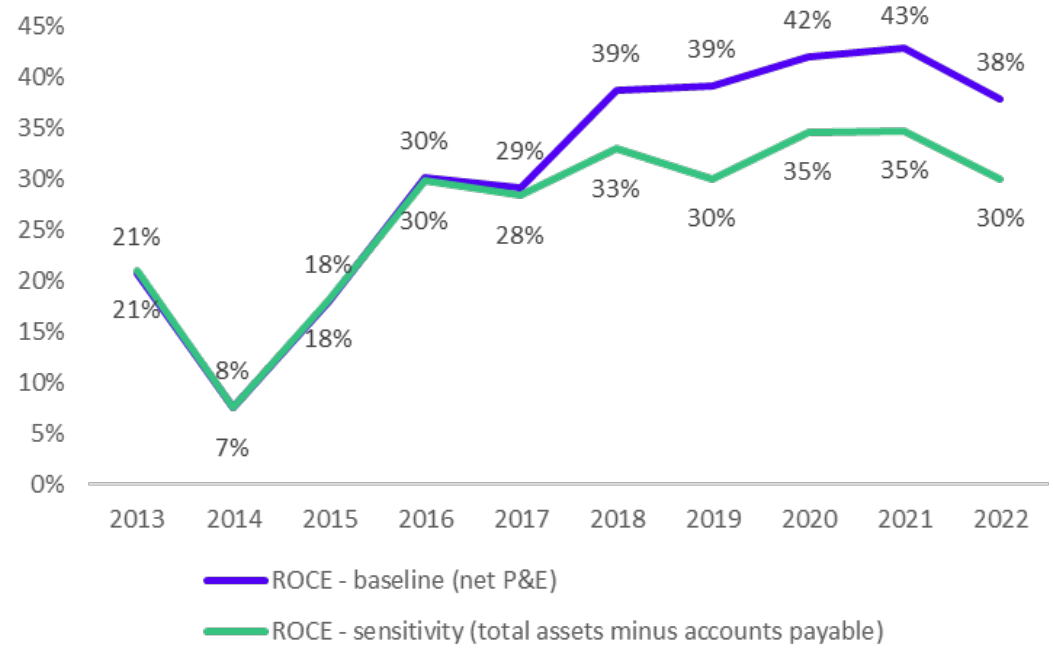

Source: Ofcom analysis based on public information from Amazon 10-K reports and Ofcom assumptions.

A2.77 Our baseline estimate indicates that AWS ROCE has increased since 2013 and averaged 40% between 2018 and 2022. It is higher than our estimate of WACC in all years except 2014, i.e. in nine of the last ten years.[116] Although not shown in Figure A2.9, we have also estimated AWS ROCE for the year to June 2023, by reference to AWS's reported EBIT in the four quarters to June 2023 and an estimate of AWS net property and equipment provided by AWS as of June 2023.[117] We estimate that AWS ROCE reduced [✂] in this period. A fall in ROCE is consistent with recent quarterly reductions in EBIT[118] but it remains significantly above WACC.

A2.78 AWS ROCE in our first sensitivity is generally lower, but also indicates that AWS ROCE has increased since 2013. Between 2016 and 2022 it was around 30% to 35%. As with our

[116] AWS made significant price reductions to some of its core compute and storage products in April 2014, which could have affected returns reported in 2014 [accessed 26 September 2023].
[117] AWS response dated 4 August 2023 to the s.174 notice dated 13 June 2023, question 14.
[118] Amazon reported that AWS EBIT fell 21% in the quarter to March 2023 (compared to the corresponding quarter in 2022) and 6% in the quarter to June 2023.

baseline ROCE, the ROCE under this sensitivity is above our estimate of WACC in all years except 2014.[119]

A2.79 Figure A2.10 shows the AWS ROCE estimates from our second sensitivity (allocation of all technology infrastructure assets to AWS) applied to both our baseline approach and first sensitivity. The data covers Amazon's financial years ending 2017 to 2022. Compared to the interim report, this chart has been updated to include 2022 data received from AWS.

A2.80 Our ROCE estimates under this sensitivity [✂].

Figure A2.10: AWS ROCE estimates (addition of non-AWS technology infrastructure assets), 2017-22

[✂]

Source: Ofcom analysis based on public information from Amazon 10-K reports and data provided by AWS.[120]

Our assessment of ROCE for Microsoft Azure

A2.81 To derive our baseline ROCE estimates for Azure, we used our estimated EBIT for Azure. We explained how we estimated this above.

A2.82 Microsoft does not separately report assets for Azure and was not able to provide us with any balance sheet information on Azure in response to our information requests.[121] We have therefore estimated capital employed for Azure as follows. This methodology is the same as that presented in the interim report, though we have updated the analysis to take account of Microsoft's 2023 financial results.

A2.83 In response to our information requests, Microsoft provided infrastructure operating costs for Azure, relating to servers and data centres, for its financial years ending 2018 to 2023. It also provided the same information for Microsoft overall.

A2.84 We also have information on Microsoft's overall property and equipment and operating lease assets. These assets will include data centres and servers used by Azure. We have estimated capital employed for Azure by calculating the ratio of these asset values to Microsoft's infrastructure operating costs, and then applying this ratio to Azure's infrastructure operating costs.

A2.85 We think this could give a reasonable estimate of capital employed for Azure as i) these assets account for around 70% of Microsoft's non-current assets over this period[122] and ii)

---

[119] One reason for the difference between our baseline ROCE and the ROCE in our first sensitivity could be that Amazon started capitalising operating leases in 2019. This could have increased the asset base and reduced ROCE from this point.

[120] AWS response dated 17 January 2023 to question 8 of our follow-up email dated 16 December 2022 concerning the s.174 notice dated 24 October 2022, Part B question 31; AWS response dated 26 May 2023 to question 6 of our follow-up email dated 5 May 2023 concerning the s.174 notice dated 24 October 2022, Part B question 31.

[121] Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, Part B questions 24 and 32; Microsoft response dated 20 January 2023 to our follow-up email dated 20 December 2022 concerning the s.174 notice dated 21 October 2022, Part B question 22; Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).

[122] After the exclusion of goodwill from non-current assets, consistent with our approach to capital employed set out above. For example, in the financial year ended June 2023, Microsoft reported net property and equipment of $95.6bn, operating lease assets of $14.3bn, total non-current assets of $227.7bn and goodwill of

Microsoft has operated a net current liability position (excluding cash and equivalents) for many years, such that the inclusion of this net current liability balance will reduce capital employed (and increase ROCE). Also, in Microsoft's last five financial years, [✂][123], indicating assets associated with these costs are likely to form a significant proportion of Azure capital employed.

A2.86 The ratio we applied to Azure's infrastructure operating costs ranges from [✂] for Microsoft's financial years ending 2018 to 2023.[124] Our estimate of Azure's ROCE is sensitive to our estimate of Azure capital employed. If the ratio of Azure's infrastructure operating cost to property and equipment and leased assets is higher than for Microsoft overall, its capital employed would be higher and its ROCE lower than our estimate.[125]

A2.87 Our ROCE estimates for Azure are shown in the figure below.

Figure A2.11: Azure ROCE estimates, Microsoft financial years ending 2018-2023

[✂]

Source: Ofcom analysis based on Microsoft 10-K reports and information provided by Microsoft in response to our information requests.[126]

A2.88 We estimate that Azure's ROCE increased over this period and was above our estimate of WACC for Microsoft's last three financial years. [✂].

A2.89 We recognise that our Azure ROCE estimates involve more assumptions than for AWS, as Microsoft reports less information than Amazon on its cloud businesses, and it could not provide us with all the financial information we requested on Azure. However, our analysis is [✂].[127]

A2.90 As with AWS, we also recognise that the replacement cost of capital employed may be different from historical values. Microsoft depreciates its servers and networking equipment over six years (four years prior to July 2022).[128] Therefore, at least for these assets, there may not be a large difference between historical value and depreciated replacement cost, given the relatively short asset lives. As with AWS, it is also possible that the replacement

---

$67.9bn. These figures are reported in page 60 of Microsoft's 2023 10-K. This means that net property and equipment plus operating leases together accounted for 69% of total non-current assets (excluding goodwill). Over the period 2018 to 2023 the ratio averaged 68%.

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

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

[125] For example, if the mix of assets associated with Azure had a longer useful life than for Microsoft overall.

[126] Microsoft response dated 20 January 2023 to our follow-up email dated 20 December 2022 concerning the s.174 notice dated 21 October 2022, Part B question 22; Microsoft response dated 7 August 2023 to the s.174 notice dated 23 May 2023, question 1b (Confidential Supplemental Annex B22).

[127] Microsoft response dated 20 January 2023 to our follow-up email dated 20 December 2022 concerning the s.174 notice dated 21 October 2022, Part B question 22; Microsoft response dated 18 November 2022 to the s.174 notice dated 21 October 2022, Part A question 33, documents: MSFT_Ofcom_Cloud_00000136, p.18; MSFT_Ofcom_Cloud_00000353, p.14; MSFT_Ofcom_Cloud_00000811, p.25; MSFT_Ofcom_Cloud_00002238, p.22; MSFT_Ofcom_Cloud_00002816, p.24.

[128] Microsoft 10-K for the year ending June 2023, page 41.

cost for some assets, such as computing, assets is lower than historical cost. For other assets there may be more of a difference between historical and replacement cost, though the margin of error would need to be relatively large for our findings to be impacted, especially in recent years. For example, we estimate that in Microsoft's 2023 financial year, [✂].

A2.91 In response to our interim report, Microsoft said that the estimation of ROCE should take account of the value of intangible assets relevant to its cloud business.[129] Microsoft did not provide examples or estimates of any intangible assets it considered relevant.

A2.92 We agree that relevant intangible assets should be included in our estimate of Azure capital employed.[130] Intangible assets can be self-created (e.g. patents, computer software) or acquired as part of business combinations.[131] A large proportion of Microsoft's reported intangibles appear to relate to acquisitions.[132] To assess if the inclusion of intangible assets could impact our ROCE findings, we considered two sensitivities.[133]

a) First, Microsoft provided us with a breakdown of all cloud-related acquisitions for its financial years ending 2018 to 2022.[134] We included an estimate of the intangible assets associated with these acquisitions in our estimate of Azure capital employed.[135] This will overestimate the impact on Azure ROCE as not all these cloud acquisitions are relevant to Azure.[136] We find that this sensitivity slightly reduces our Azure ROCE but does not impact our finding that Azure ROCE was above our estimate of WACC for Microsoft's last three financial years.
b) Second, we capitalised an estimate of Azure R&D operating expenses,[137] assuming a useful life of between 3 and 5 years.[138] We find that our ROCE estimates for Azure in its 2022 and 2023 financial year increase by up to 2 percentage points when we capitalise

[129] Microsoft response to the interim report, page 38, paragraph 150.
[130] The CMA's Guidelines for market investigations set out criteria for the inclusion of intangible assets when estimating capital employed. See Competition Commission, 2013. Guidelines for market investigations: Their role, procedures, assessment and remedies, Annex A paragraph 14 [accessed 26 September 2023]. The criteria are that it must comprise a cost that has been incurred primarily to obtain earnings in the future, ii) the cost must be additional to costs necessarily incurred at the time in running the business and iii) it must be identifiable as creating such an asset separate from any arising from the general running of the business.
[131] Under accounting standards, intangible assets are recognised when they are identifiable, the cost can be measured reliably and where they will give rise to future economic benefits.
[132] For example, Microsoft's 2023 10-K reports gross intangible assets of around $23.5bn ($9.4bn net), the majority of which appear to relate to intangible assets associated with acquisitions in recent years, e.g. LinkedIn (2016), GitHub (2018) and Nuance Communications (2022).
[133] These sensitivities were carried out to understand if including intangible assets potentially associated with Azure could affect our findings. We have not carried out an exercise to assess if these assets would meet the requirements set out in the CMA Guidelines for inclusion as an intangible asset.
[134] Microsoft response dated 9 December 2022 to the s.174 notice dated 21 October 2022, Part B question 27 (Confidential Annex B27).
[135] We also made a similar adjustment to Azure capital employed for Microsoft's financial year ending 2023 based on intangible assets data reported in pages 78-79 of Microsoft's 2023 10-K.
[136] We estimate that only a small proportion of these cloud acquisitions are relevant to Azure. For example, the two largest cloud acquisitions in the data provided by Microsoft relate to Nuance Communications and GitHub – neither of which Microsoft reports as part of Azure.
[137] Microsoft could not provide us with R&D data at the Azure level, as noted above, so we estimated this based on information it provided on R&D at the Microsoft and Cloud and Enterprise level.
[138] This useful life assumption appears broadly consistent with Microsoft's assigned useful life for technology-based intangible assets. In recent years, where Microsoft's 10-K reports its useful life assumption for technology-based intangible assets, it has typically assumed 3 to 5 years.

Azure R&D. This is because capitalising Azure R&D increases Azure EBIT by relatively more than capital employed, which overall slightly increases ROCE.

A2.93 Overall, these sensitivity analyses suggest that adding intangible assets to Azure capital employed is unlikely to change our conclusions regarding the trajectory and level of Azure ROCE.

Our assessment of ROCE for Microsoft Cloud

A2.94 To derive our baseline ROCE estimates for Microsoft Cloud, we used our estimated EBIT for Microsoft Cloud. We explained how we estimated this above.

A2.95 Microsoft does not separately report assets for Microsoft Cloud.[139] As a result we estimated capital employed for Microsoft Cloud by reference to capital employed for Microsoft's overall business. This approach is the same as that taken in the interim report, though we have updated the analysis for Microsoft's 2023 financial results. This is a conservative approach which will overestimate capital employed (and underestimate ROCE) given that Microsoft Cloud represented a third of Microsoft's revenue, on average, in its financial years ending 2016 to 2023.

A2.96 Our baseline estimate of capital employed (and ROCE) is calculated by reference to the net property and equipment assets plus operating leases for Microsoft's entire business. This is because these types of assets are likely to be relevant to providing Microsoft Cloud's services. This approach broadly aligns with our baseline approach to calculating AWS ROCE above.

A2.97 As a sensitivity, we consider the impact of additionally including Microsoft's published net accounts receivable position (i.e. the difference between Microsoft's accounts receivable and accounts payable) as part of capital employed. This is comparable to our first sensitivity analysis for the AWS ROCE, which subtracts an AWS accounts payable estimate from total AWS assets (which includes AWS accounts receivable).[140]

A2.98 The chart below shows our baseline ROCE estimates for Microsoft Cloud, alongside the estimate from our sensitivity, for Microsoft's financial years ending 2016 to 2023 (Microsoft's financial year ends in June).

[139] Microsoft's 10-K reporting says that assets are not allocated to segments for internal reporting presentations (see page 95 of its 2023 10-K). In response to our information requests, Microsoft was unable to provide capital employed data for Microsoft Cloud – Microsoft responses dated 31 May 2023 and 7 August 2023 to the s.174 notice dated 23 May 2023, question 1a and 1b.

[140] It is possible that there are other current and non-current assets and current liabilities from Microsoft's balance sheet which are relevant to the capital employed for Microsoft Cloud. However, over the period considered, the sum of potentially relevant current liabilities (total current liabilities minus accounts payable, which is already included in our sensitivity analysis) exceeds potentially relevant current and non-current assets (inventories, other current assets, equity investments, net intangible assets and other long-term assets). This equates to a net liability position, which, if included, would reduce total capital employed and increase ROCE. We have therefore not tested the ROCE impact of including these other assets and liabilities.

Figure A2.12: Baseline ROCE and sensitivity ROCE estimates for Microsoft Cloud, Microsoft financial years ending 2016-2023

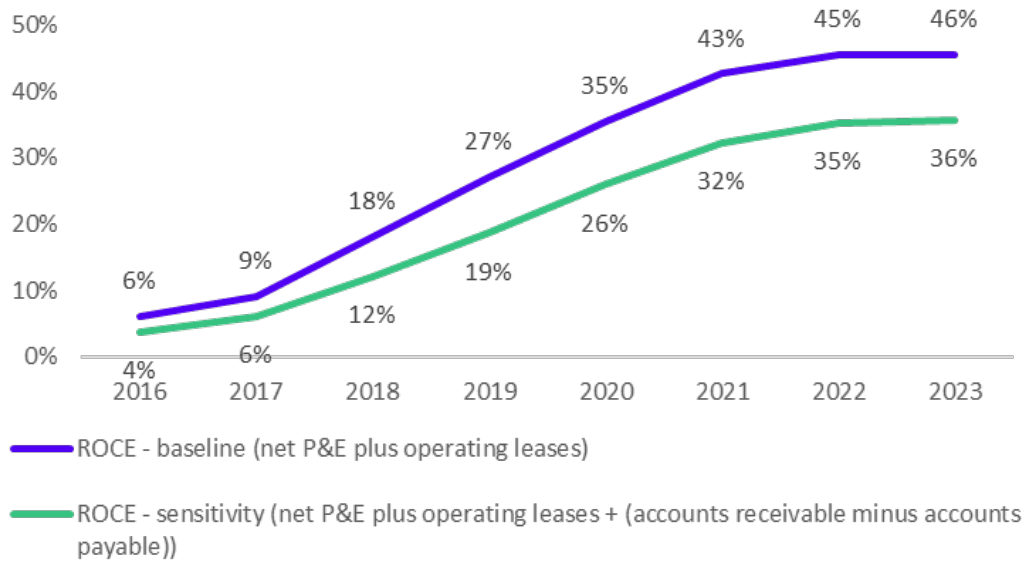

Source: Ofcom analysis based on public information from Microsoft 10-K reports and Ofcom assumptions.

A2.99 We estimate that Microsoft Cloud ROCE steadily increased over this period. In our baseline, conservative estimate of capital employed, we estimate Microsoft Cloud ROCE ranged from 18% to 46% in Microsoft's financial years ending 2018 to 2023, above our estimate of WACC.

A2.100 Adding Microsoft's net accounts receivable position to capital employed in our sensitivity reduces 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.

Relevance and interpretation of our findings

A2.101 As noted above, profitability can be one indicator of how well competition is working and we consider it alongside other indicators in Section 8 as part of our overall assessment of competition in cloud infrastructure.

A2.102 Our ROCE analysis above indicates that:

a) AWS returns have been significantly above WACC in every year since 2014 and between 2018 and 2022 ROCE ranged from 38% to 43%, averaging 40%.
b) Azure's returns increased between Microsoft's 2018 and 2023 financial years and were above WACC in Microsoft's last three financial years. Microsoft Cloud returns were above WACC between Microsoft's 2018 and 2023 financial years.

A2.103 We also noted that EBIT margins for AWS and Microsoft Cloud have been higher than other cloud providers, for whom operating losses have been common. We said Azure operating profits [✂].

A2.104 We received comments from stakeholders on the relevance of our profitability analysis to the market study and how analysis should be interpreted. These are summarised below alongside our response.

Comments on the relevance of our profitability analysis

A2.105 Microsoft said comparing ROCE to WACC is "*not informative as to competitiveness in the cloud services market*".[141] Microsoft said that while WACC can be used as a comparative benchmark for ROCE, the expectation that ROCE should converge to equal WACC over the period considered is misguided. Microsoft said that returns can and should consistently exceed the cost of capital for pro-competitive reasons such as successful innovation and operating efficiency.[142]

A2.106 Microsoft also said that considering ROCE over a one-year or five-year period will not capture the dynamic nature of cloud investment profitability (for example, because large upfront investments are required which generate returns over a long time period) and will ignore early periods where its cloud business was significantly less profitable.[143]

A2.107 We disagree that comparing ROCE to WACC is not informative. As set out above, this is a standard approach to assessing firms' profitability and consistent with CMA Guidelines for market investigations.[144] We agree that returns can exceed WACC for a number of reasons such as successful innovation and greater efficiency, but in a competitive market we would expect to see returns tending towards the WACC over time (though not necessarily to equal WACC) as competitors enter the market and compete. Where high returns are observed alongside barriers to switching and multi-homing, as is the case in cloud infrastructure services, returns persistently exceeding WACC could indicate the limitations to effective competition.

A2.108 We agree that analysing returns over a multi-year period is important as we are interested in understanding whether returns have been persistently higher than WACC, which could indicate limitations to the competitive process. For AWS we were able to analyse 10+ years of data, but for Azure we only had 6 years of data. We think this is sufficient to understand how Azure's returns have developed and the extent to which they have been above WACC. Extending our analysis of Azure to cover more years would not change our high-level findings that Azure returns have been above WACC in Microsoft's last three financial years, but we agree that in earlier periods Azure is likely to have been significantly less profitable.

A2.109 AWS made two comments on profit margins. First, it said it was hard to tell if profit margins are higher than might be expected when competition is working well as there was no competitive margin against which to compare them.[145] Second, it said there is a wide dispersion in profit margins across the range of services provided by cloud services providers and suggested this means analysing global average margins is unlikely to provide robust or meaningful results.[146]

A2.110 We agree there is no clear benchmark against which to compare EBIT margins for the purposes of assessing profitability in the context of competition analysis. As a result, EBIT margins on their own do not indicate whether returns are higher than might be expected in a market that is working well. The purpose of our EBIT analysis was to compare EBIT margins

---

[141] Microsoft response to the interim report, page 39, paragraph 153.
[142] Microsoft response to the interim report, page 39, paragraph 152.
[143] Microsoft response to the interim report, pages 38-39, paragraph 151.
[144] Competition Commission, 2013. Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraphs 114-125 [accessed 26 September 2023].
[145] [✂].
[146] [✂].

across cloud providers as well as serving as an input to our ROCE analysis, on which we place more weight as there is a competitive benchmark against which to compare it (the WACC, as set out above). The fact that AWS's EBIT margins have been significantly higher than other global cloud providers, many of which have not been profitable to date, could indicate it is difficult for competitors to sustainably enter the cloud infrastructure services market. This is likely to reflect the barriers to entry and expansion we identify in Section 6.

A2.111 We disagree that analysis of global average EBIT margins across AWS services (and by extension ROCE, which uses these margins) is uninformative. Capital employed (e.g. investment in datacentres and servers) supports the provision of all cloud infrastructure services, so to understand returns on capital employed for cloud infrastructure activities we consider it is appropriate to take account of AWS's overall profitability across all its services.

Comments on the interpretation of our profitability analysis

A2.112 Based on our ROCE analysis, it appears that AWS returns have been persistently above WACC since 2014 while Azure returns increased between Microsoft's 2018 and 2023 financial years and were above WACC in Microsoft's last three financial years. Given the presence of barriers to multi-cloud and switching, alongside the persistence of these profits even though there has been entry (e.g. by Oracle in 2016), these findings indicate there could be limitations in the competitive process.

A2.113 Some responses to our interim report said other factors could also explain returns above WACC. AWS and Microsoft both suggested that high returns are expected and a necessary incentive where, as is the case with cloud infrastructure services, firms innovate and make high upfront investments.[147] Microsoft also said returns can exceed the WACC for reasons such as operating efficiency and that publicly traded companies in competitive industries frequently have ROCE greater than WACC over sustained periods.[148]

A2.114 We agree in principle that these factors (i.e. innovation combined with high upfront investments, and operating efficiency) could explain some of the returns above WACC, at least for a period. Temporary high profits are common in well-functioning markets and provide rewards for developing comparative advantage as well as giving signals and incentives for other firms to invest. However, as set out above, in a market characterised by effective competition, we would not expect returns to be persistently above WACC as competitors would enter and compete away the difference.

A2.115 We explain below why, based on the evidence available, we do not consider that successful innovation combined with high upfront investment can fully explain the magnitude and persistency of returns above WACC that we estimate, especially for AWS.

A2.116 In respect of operating efficiency, Microsoft did not assert or attempt to demonstrate that its returns were in fact due to greater operating efficiency, it simply raised the possibility that they could be.

A2.117 The possibility that its returns could be explained, in part, by greater operating efficiency does not exclude the possibility of competition concerns, particularly where the source of

[147] AWS said high returns can be a feature of industries where competition is driven by constant innovation and firms make high upfront investments in assets, and Microsoft said our analysis should take account of high levels of investment and that returns can exceed the WACC due to successful innovation. [✂]; Microsoft response to the interim report, pages 38-39, paragraphs 150-152.
[148] Microsoft response to the interim report, page 39, paragraph 152.

some cost advantages may itself reflect barriers to entry and expansion. These could include some of the features we identified in Section 6 that could drive lower unit costs, such as economies of scale in data centres (due to larger data centres, discounts on hardware like servers, and higher server utilisation), innovation in hardware like custom processors, and experience curve effects associated with access to skilled staff and earlier entry.

A2.118 In a competitive market, we would expect that returns above the WACC would reduce over time, as new entrants enter the market and become more efficient as they gain scale. As we explain in Section 8, the barriers to switching and multi-cloud may inhibit the ability of smaller cloud providers to gain scale and achieve the same cost efficiency as the market leaders, and in turn are less able to put pressure on the high returns earned by AWS and Microsoft. We recognise that, in the case of Microsoft Azure it may be too early to determine whether the current level of profits will persist, but this does seem to be the case for AWS based on our analysis set out above.

A2.119 The magnitude of returns above the WACC and, at least in the case of AWS, their persistence, is therefore consistent with an assessment that a market that is not working well.

Innovation combined with high upfront investments

A2.120 Returns above WACC could be consistent with an outcome where actual returns are above ex-ante expected returns purely because a risky investment has succeeded. For example, the possibility of making high returns if risky innovation succeeds can be necessary for companies to take the risk of innovating in the first place. For a risky project to have expected returns in line with the WACC, there needs to be a sufficient chance of making returns above the WACC if the innovation succeeds, to offset the risk that returns will be less than the WACC if it fails. Actual returns above the WACC could therefore be the consequence of successful innovation and, if not sustained over the longer term, be consistent with a well-functioning competitive process.

A2.121 We consider the argument that high returns are a consequence of innovation and high upfront investment is primarily applicable to business activities that require substantial investment in non-scalable activities with significant demand risk, and where expected returns are not above the WACC at the start of the investment but actual returns could be if the innovation succeeds.

A2.122 Companies with scalable investments would not be expected to experience returns persistently and significantly above WACC because capital is not at risk for long periods of time and the impact of demand shocks would be lower, i.e. the company could scale back its investments in response. Similarly, where large non-scalable investments are made against high levels of demand risk, the more likely it is that high ex-post returns could be observed compared to ex-ante expected returns. Therefore, where investments are scalable and demand uncertainty is relatively low, returns persistently and significantly above WACC would not be required to compensate for downside risk.

A2.123 Based on the evidence available, we think it is unlikely these factors are a significant feature of cloud infrastructure activities that could fully explain returns above WACC, especially the sustained returns significantly above WACC we estimate for AWS, for the following reasons.

a) **Some cloud investments appear to be relatively scalable**. While AWS has made significant investments over time,[149] Figure A2.8 illustrates that its asset base has increased from a relatively low level, and broadly increased in line with revenue. In 2013, AWS capital expenditure was $2bn; while in 2022 it was almost $28bn, indicating that in the early years of AWS, even if investments were less scalable, or demand was more uncertain than today, less capital was at risk. Amazon's 2021 letter to shareholders provides support for initial investments being relatively small: "*We launched EC2 in 2006 with one instance size, in one data center, in one region of the world, with Linux operating system instances only (no Windows), without monitoring, load balancing, auto-scaling, or yes, persistent storage. EC2 was an initial success, but nowhere near the multi-billion-dollar service it's become until we added the missing capabilities listed above, and then some*".[150]

b) **Some cloud investments appear to be informed by demand signals**. Recent comments from cloud providers suggest some investments are made once demand is established. For example, on a recent quarterly earnings call Microsoft said its "*datacentre investments continue to be based on strong customer demand and usage signals*"[151] and Oracle said *"we remain careful to pace our investments appropriately and in line with booking trends, which is why our gross margins are up in our cloud"*.[152] AWS also noted that some of its services were originally developed in response to a single customer's needs, before being rolled out more widely to benefit other customers, which might suggest some investments are only made once demand is more certain.[153] Further, we note in Section 5 that hyperscalers said that customer commitments allow them to forecast future demand more accurately, allowing them to invest more confidently, which would also be consistent with investments being informed by demand signals. It is possible that in earlier years, demand risk associated with cloud investments was higher, though AWS revenue growth since 2013 does not appear particularly volatile compared to other S&P 500 companies and volatility appears to have decreased in recent years.[154]

A2.124 This evidence suggests that returns persistently above the WACC were not needed to compensate for downside risks associated with risky innovation or large upfront investments. There is also some evidence from AWS that expected returns were above the WACC some years ago, which would indicate that high actual returns were not driven by successful risk-taking. For example, Amazon's 2014 letter to shareholders says: "*Finally, I'm optimistic that AWS will have strong returns on capital. This is one we as a team examine because AWS is capital intensive. The good news is we like what we see when we do these analyses. Structurally, AWS is far less capital intensive than the mode it's replacing – do-it-yourself datacenters – which have low utilization rates, almost always below 20%. Pooling of workloads across customers gives AWS much higher utilization rates, and correspondingly higher capital efficiency. Further, once again our leadership position helps: scale economies*

[149] For example, AWS told us it incurred around [✂] of gross data centre capital expenditure globally between 2019 and 2021. Source: AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, Part B question 26. In March 2022, AWS also announced that it would invest £1.8bn over the next two years in building and operating data centres in the UK [accessed 26 September 2023].
[150] Amazon 2021 letter to shareholders.
[151] Microsoft, Q1 FY2023 earning calls transcript.
[152] The Motley Fool, 2023. Oracle (ORCL) Q4 2023 Earnings Call Transcript [accessed 26 September 2023].
[153] [✂].
[154] Measured using the standard deviation of annual revenue growth rates between 2013 and 2022.

can provide us a relative advantage on capital efficiency. We'll continue to watch and shape the business for good returns on capital".[155]

[155] Amazon, 2014 letter to shareholders. We asked Amazon for the analysis supporting its 2014 view that AWS could have strong returns on capital, [✂] (AWS response dated 9 December 2022 to the s.174 notice dated 24 October 2022, Part B question 34).

A3. Prevalence of multi-cloud and switching

A3.1 In this annex, we set out evidence on the prevalence of multi-cloud and switching. As part of this, we present findings from the Context Consulting market research and other industry reports (which we originally included in Annex 7 of our interim report), drawing out differences and explanations for those differences, where we identify them. In addition, we summarise the positions on multi-cloud and switching presented by AWS and Microsoft in response to our interim report and present our assessment of their positions. We also review the quantitative evidence on multi-cloud, switching and win rates submitted by [✂].

A3.2 As set out in Section 3, it is possible to broadly categorise multi-cloud in several different ways. The categorisation we have adopted (notably: cloud duplication, siloed multi-cloud and integrated multi-cloud) has been developed based on the evidence we have gathered and assessed during the market study and we consider it useful for our purpose of assessing the state of competition in cloud infrastructure services.

A3.3 We commissioned market research conducted by Context Consulting at an early stage of the market study process, before our categorisation of the different forms of multi-cloud presented in Section 3 was fully fleshed out. For this reason, in some cases, the findings from the Context Consulting market research may not exactly map onto one of the specific forms of multi-cloud that we have adopted.

A3.4 We also note that the Context Consulting market research asked respondents about multi-cloud and switching in relation to IaaS/PaaS providers. Customers may have understood such providers to include both cloud providers and ISVs. This means that responses regarding the use of multi-cloud are likely to include cases where a customer is using first-party and third-party services on the same cloud (i.e. in what we refer to as multi-vendor architectures earlier in this report) or combining public and private cloud (i.e. hybrid cloud). Similarly, responses in relation to switching are likely to cover switching between different public clouds, switching within the same public cloud, or switching between on-premises IT/private cloud and public cloud.

A3.5 To capture a richer and more in-depth understanding of customer responses on multi-cloud and switching, we also commissioned qualitative market research from Context Consulting. Overall, we consider that the quantitative and qualitative stages of the Context Consulting market research provide useful evidence on multi-cloud and switching, albeit with some caveats (as we discuss below).

Evidence on multi-cloud

Prevalence of multi-cloud

A3.6 About 52% of IaaS and/or PaaS users in the market research use more than one IaaS/PaaS provider. Smaller businesses, technology “laggards”, and public sector organisations are more likely to use just one IaaS/PaaS provider. If we exclude respondents who said they were only using private cloud, the proportion of respondents who were using more than one IaaS/PaaS provider increases to 61%.

A3.7 We note that these figures are likely to overstate the fraction of customers using multiple public cloud providers (i.e. 'multi-cloud' as defined in this report). This is because, as noted above, customers may have understood IaaS/PaaS providers to include private cloud providers, public cloud providers and ISVs. Hence, in addition to customers using multiple public IaaS/PaaS cloud providers, the 52% figure above may capture:

a) customers using cloud providers and ISVs hosted on the same cloud (i.e. multi-vendor architectures);
b) customers using a combination of private and public cloud providers (i.e. hybrid cloud); or
c) customers using multiple private cloud providers. The 61% figure above would exclude group c), but may still include customers from groups a) and b).

A3.8 The market research also asked respondents who use only one IaaS/PaaS provider and those considering purchasing IaaS/PaaS[156] about their attitude to **potentially using multi-cloud architecture in the future**. A majority responded, "yes definitely" (23%) or "yes possibly" (63%).[157]

A3.9 The prevalence of multi-cloud use in the Context Consulting market research is much lower if we compare it with some industry reports. For example, the Flexera 2023 State of the Cloud report suggests that 87% of companies and organisations they surveyed were using multi-cloud (out of which: 72% were hybrid cloud; 2% - multiple private; and 13% - multiple public).[158] This may be explained by the fact that the Context Consulting market research was aimed at IaaS and/or PaaS use only. The Flexera report is also more skewed towards larger organisations (83% of their respondents have more than 1,000 employees compared to only 29% in our sample) and has a wider geographical representation whereas the Context Consulting market research concentrated on the UK.

A3.10 A 2023 report commissioned by Oracle suggests that 98% of respondents were using (or planning to use within 6 months) more than one IaaS/PaaS provider.[159]

A3.11 Foundry's Cloud Computing Study 2022 reports that in their sample 16% of organisations rely on a single cloud provider for their public cloud deployments (19% for small and medium businesses and 13% for "enterprises", i.e. respondents with more than 1,000 employees).[160]

A3.12 Public First conducted a survey among 1,001 current or potential users of IaaS/PaaS/SaaS services in the UK, including 716 IaaS/PaaS users.[161]

A3.13 Of these 716 IaaS/PaaS users in Public First's survey, 41% (i.e. 295 respondents) are from large companies, i.e. with more than 500 employees. In addition, the majority of these 716 IaaS/PaaS users come from companies in the ICT sector (22%, i.e. 161 respondents), followed by the financial and insurance sector (12%, i.e. 83 respondents), manufacturing

---

[156] Those who did not use IaaS and/or PaaS at the time of the market research but were actively considering using IaaS and/or PaaS within 6 months.
[157] Context Consulting market research report, slide 79.
[158] Flexera, 2023. 2023 State of the Cloud report, page 18 [accessed 22 August 2023].
[159] Multicloud in the Mainstream: Making IT Work 'As Advertised'. Commissioned by Oracle. S&P Global Market Intelligence. February 2023. Figure 1, p.1.
[160] https://resources.foundryco.com/download/cloud-computing-executive-summary, page 3 [accessed 22 August 2023].
[161] The results of the Public First's survey are available at: https://www.publicfirst.co.uk/files/CCIA_Survey.xlsx [accessed 7 September 2023].

(11%, i.e. 77 respondents) and wholesale and retail trade (11%, i.e. 76 respondents). Furthermore, a relatively high proportion of the 716 IaaS/PaaS users (32%, i.e. 230 respondents) uses private cloud only, compared to 17% (i.e. 124 respondents) who uses public cloud only. Lastly, slightly less than 50% of the 716 IaaS/PaaS users use both private and public clouds, with 39% (i.e. 277 respondents) saying that private and public clouds are integrated with each other and 9% (i.e. 67 respondents) saying that the two clouds are part of separate IT architectures.[162]

A3.14 According to the results of the Public First's survey, 71% of IaaS/PaaS users (i.e. 509 out of 716 respondents) said that they use more than one cloud infrastructure provider.[163] As is the case for the Context Consulting research, this is likely to capture a range of different cloud architectures which do not necessarily involve the use of two public cloud providers.[164]

A3.15 In addition, 56% of IaaS/PaaS users (i.e. 401 out of 716 respondents) said that they would be likely to add an additional cloud infrastructure provider in the next few years, compared to 22% (i.e. out of 716 respondents) who said they were neither likely nor unlikely to do so and 18% (i.e. out of 716 respondents) who said they were unlikely to do so.[165]

A3.16 In response to our interim report,[166] AWS said that multi-cloud is a popular approach and well used; and that the Context Consulting market research commissioned by Ofcom also supports this view. However, according to AWS, other public reports (already mentioned in earlier sections of this annex)[167] suggest that the findings of the Context Consulting market research might underestimate the actual popularity of multi-cloud.

---

[162] See question 18 of the Public First's survey [accessed 7 September 2023].

[163] See question 19 of the Public First's survey [accessed 7 September 2023].
The main reason for using more than one cloud provider was "Increased flexibility", followed (in order of preference) by "Performance optimisation", "Risk mitigation", "Cost optimisation" and "Being able to use new or best technologies or features from multiple providers". Slightly less frequently mentioned reasons were "Compliance or data sovereignty requirements", "To make sure we're not locked into one supplier" and "Would allow you to switch more easily if needed". See question 40 of the Public First's survey [accessed 7 September 2023].

[164] As in the case of our Context Consulting market research, we cannot rule out that, in addition to cases of customers using multiple public IaaS/PaaS cloud providers (i.e. the definition of multi-cloud used in this report), the Public First's results also captures cases of customers using a cloud provider and an ISV hosted on the same cloud (i.e. multi-vendor architectures), customers using more than one supplier of cloud services across public and private cloud (i.e. hybrid cloud), or customers using multiple suppliers of private cloud.

[165] See question 56 of the Public First's survey [accessed 7 September 2023].
Amongst those unlikely to add an additional cloud provider in the next few years, the most cited reasons (in order of preference) were: i) Satisfied with features and value for money of current providers; ii) Have no unmet cloud needs; iii) Would be too expensive to add another provider; and, iv) Would be too complicated to add another provider. See question 57 of the Public First's survey [accessed 7 September 2023].

[166] AWS response to the interim report, paragraph 10.

[167] Notably: i) the Flexera 2023 State of the Cloud report (which found that 87% of respondents use multiple clouds; ii) the 2023 report commissioned by Oracle (which found that 98% of respondents use or plan to use more than one cloud provider; and iii) the Gartner 2021 report (which indicated that 78% of customers in the UK worked with different cloud infrastructure providers). Microsoft also referred to the findings of the Context Consulting market research commissioned by Ofcom as well as to the third-party surveys mentioned earlier. In addition, Microsoft listed several UK and European companies "with customer stories featured by more than one provider" as examples of companies using more than one cloud provider. Microsoft response to the interim report, Annex 5.

Management of multi-cloud architecture

A3.17 The market research asked respondents who use more than one IaaS/PaaS provider about the way they manage their multi-cloud architecture. The largest proportion (45%) use different IaaS/PaaS providers for different workloads. Spreading similar workloads across IaaS/PaaS providers is the second most frequent response (at 40% overall) and 15% have one main IaaS/PaaS provider and use others as a back-up.[168]

A3.18 As set out above, it is not always feasible to map the market research results onto one of the specific forms of multi-cloud we have set out in Section 3. For example, the first result (i.e. customers using different IaaS/PaaS providers for different workloads) and the third result (i.e. customers who have one main IaaS/PaaS provider and use others as a back-up) may be indicative of siloed multi-cloud and cloud duplication respectively. However, the second result (i.e. customers spreading similar workloads across IaaS/PaaS providers) may be indicative of both siloed and integrated multi-cloud.[169]

A3.19 In addition, as noted above, we consider the reported frequencies are likely to overstate prevalence of multi-cloud more generally. This is because customers may have understood IaaS/PaaS providers to include private cloud providers, public cloud providers and ISVs. This means that, in addition to cases of customers using multiple public IaaS/PaaS cloud providers (i.e. the definition of multi-cloud used in this report), these results may capture cases of customers using a cloud provider and an ISV hosted on the same cloud (i.e. multi-vendor architectures), customers using more than one supplier of cloud services across public and private cloud (i.e. hybrid cloud), or customers using multiple suppliers of private cloud.

A3.20 Our market research also asked what actions customers are using to mitigate risks of lock-in and found that 25% of IaaS/PaaS users are 'using a multi-cloud strategy where workloads of an individual app are run in more than one cloud and integrated'. We note that this was the only 'multi-cloud' option in response to this question. As such, while this may capture customers who are using multiple public clouds to mitigate risks of lock-in,[170] we consider this is likely to capture all types of multi-cloud (i.e. including also siloed multi-cloud and cloud duplication). Indeed, the sample of users that selected this option can be broken down as follows: 47% are customers who use different IaaS/PaaS providers for different workloads, 39% are customers who are spreading similar workloads across IaaS/PaaS providers and 14% are customers who have one main IaaS/PaaS provider and use others as a back-up. These frequencies are very similar to the ones found across all respondents.

A3.21 According to the results of the Public First's survey, amongst the 509 (out of a total of 716) IaaS/PaaS users who use more than one cloud infrastructure provider, the majority (52%, i.e.

---

[168] Context Consulting market research report, slide 76.

[169] It is difficult to assign these responses to a specific multi-cloud architecture. For example, this may reflect customers who are: (i) running similar, but distinct, applications on separate clouds (which would be closer to a siloed multi-cloud approach); or (ii) using different clouds across geographies; or (iii) integrating services from multiple providers.

[170] We consider this figure is likely to overstate prevalence of multi-cloud. This is because, as noted above, customers may have understood IaaS/PaaS providers to include private cloud providers, public cloud providers and ISVs. This means that, in addition to cases of customers using multiple public cloud providers (i.e. the definition of multi-cloud used in this report), this result may capture cases of customers using a cloud provider and an ISV hosted on the same cloud (i.e. multi-vendor architecture), customers using more than one supplier of cloud services across public and private cloud (i.e. hybrid cloud), or customers using multiple suppliers of private cloud.

265 out of 509 respondents) said that they use different cloud infrastructure providers for a mix of integrated and independent uses. In addition, 37% of these IaaS/PaaS users (i.e. 188 out of 509 respondents) said that their cloud infrastructure providers largely work independently and in parallel together. Lastly, only 10% (i.e. 51 out of 509 respondents) said that their cloud infrastructure providers are largely integrated together.[171] It would therefore appear that, according to the Public First's survey results, integrated multi-cloud is not used extensively.[172]

A3.22 In the course of this market study, and after we commissioned the Context Consulting market research, we issued a customer questionnaire to several companies, asking them about their multi-cloud use. We note that most of the companies who responded to our request for information are relatively large companies and, therefore, not necessarily representative of the universe of cloud customers in the UK.

A3.23 While some respondents used several cloud providers, very few told us that they were using integrated multi-cloud. [✂] told us they use this approach in a very limited number of cases, and it necessarily includes additional complexity in the workload design and reduces optimisation in use of the underlying cloud services.[173]

A3.24 Another customer, [✂], explained it 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.[174] It noted that it would be cost prohibitive to maintain both clouds actively all the time, and therefore the back-up cloud of Google is more basic, but would enable the firm to scale up at speed if needed.

A3.25 Several other companies told us that they were using different cloud providers for separate workloads. For example, [✂] explained that they use AWS for their digital services (website, app, backend etc.) and Google for big data and analytics.[175] Similarly, [✂] uses a "public cloud deployment model primarily composed of IaaS products from AWS and a small public cloud deployment with Google Cloud Platform (GCP) for specialized use cases" (e.g. Google Analytics).[176]

---

[171] See question 49 of the Public First's survey [accessed 7 September 2023].
As noted in paragraph A3.13 above, amongst the 716 IaaS/PaaS respondents included in the Public First's survey, 32% uses private cloud only; 17% uses public cloud only; and 39% uses both public and private clouds. As in the case of our Context Consulting market research, we cannot rule out that, in addition to cases of customers using multiple public IaaS/PaaS cloud providers (i.e. the definition of multi-cloud used in this report), the Public First's results also captures cases of customers using a cloud provider and an ISV hosted on the same cloud (i.e. multi-vendor architectures), customers using more than one supplier of cloud services across public and private cloud (i.e. hybrid cloud), or customers using multiple suppliers of private cloud.
[172] When asked what type of integration they considered to be important, respondents selected among the following four options to a similar extent: i) Application integration (e.g. applications are able to communicate and interact with each other); ii) Management integration (e.g. having the ability to manage your different cloud platforms from a single interface); iii) Security integration (e.g. having the ability to coordinate security policies and controls across platforms); and, iv) Data integration (e.g. data synchronisation, data management, data transformation etc).
See questions 50 and 66 of the Public First's survey [accessed 7 September 2023].
[173] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]. Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email dated [✂].
[174] [✂].
[175] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].
[176] [✂] response dated [✂] to the s.174 notice dated [✂] question [✂].

A3.26 One customer [✂] told us that [✂] is its primary cloud provider but it uses Google for data analytics and it has developed some bespoke extensions on Oracle Cloud, using some of Oracle's PaaS tools.[177] It explained that in general, its preference is to purchase all cloud services from a specific cloud provider ([✂]) unless there is an operational reason not to.[178] This customer suggested there are significant advantages from having all their cloud services with a single cloud provider:

a) The ability to devote resources to recruiting, training and retraining a dedicated team, specialising in architecture of just one cloud provider.
b) The integration cost and effort is reduced and there are administrative advantages from having a single billing system.
c) The ability to qualify for discounts or, more specifically, to meet usage commitments.[179]

A3.27 Some of the larger customers who are already multi-clouding ([✂]) commented that when a new workload appears, they consider to which of their several cloud providers they could best allocate it.[180] Even in these two cases, however, the resulting cloud architecture is more oriented towards siloed multi-cloud than integrated multi-cloud.

A3.28 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.[181]

A3.29 Our findings regarding management or use of multi-cloud are generally in line with industry reports. The Flexera 2023 State of the Cloud report suggests that while organisations are using multiple clouds, this does not always mean that individual applications are spanning clouds. "Apps siloed on different clouds" is the most common multi-cloud implementation, with 44% of respondents calling it the most popular architecture for the third year in a row. "DR/Failover"[182] is now at 42%.[183]

[177] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].
[178] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[179] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[180] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[181] Context Consulting market research report, slide 80.
[182] DR is disaster recovery.
[183] Flexera, 2023. 2023 State of the Cloud report, page 20 [accessed 22 August 2023].

Figure A3.1: Use of multi-cloud in the Flexera 2023 State of the Cloud report[184]

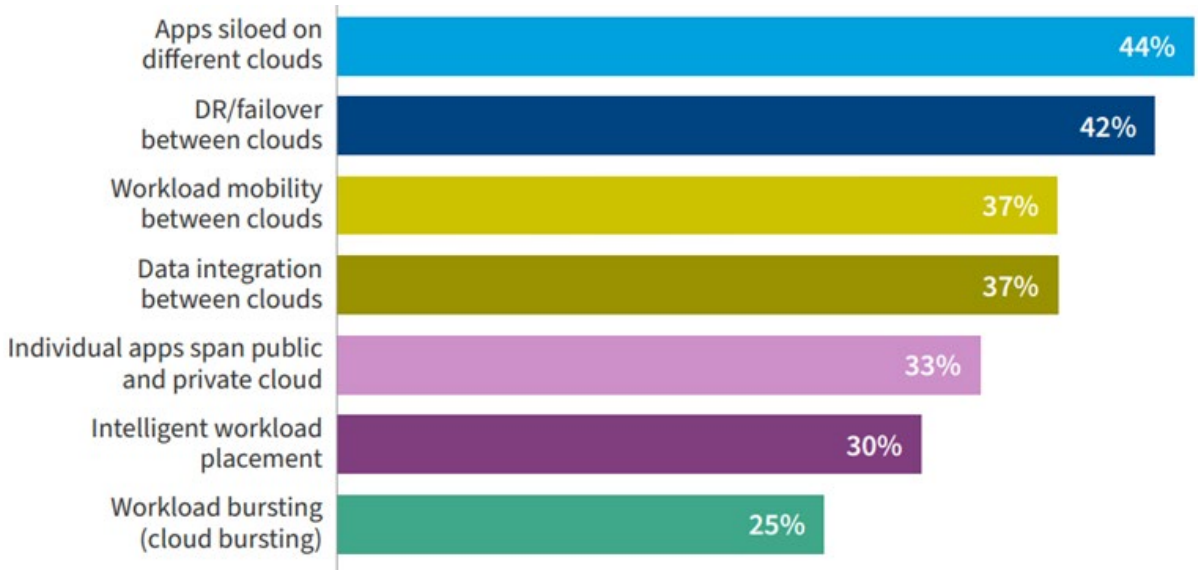

Source: Flexera 2023, State of the Cloud report, page 20.

Benefits and challenges of using multi-cloud

A3.30 All respondents to the market research (those already using more than one IaaS/PaaS provider, those using just one, and those considering purchasing IaaS/PaaS) were asked about their perceived **benefits of using multi-cloud architecture**. Overall, 48% cite additional resilience in the event of outages, 48% indicated less dependence on one supplier, 43% say it would allow them to switch more easily if needed, and 48% say it allows them to use best services/products from different suppliers.[185] Figure A3.2 below shows that respondents who already use more than one IaaS/PaaS provider tend to report benefits more often than those who only use one provider.

Figure A3.2. Benefits of using multi-cloud, by the number of providers respondents use

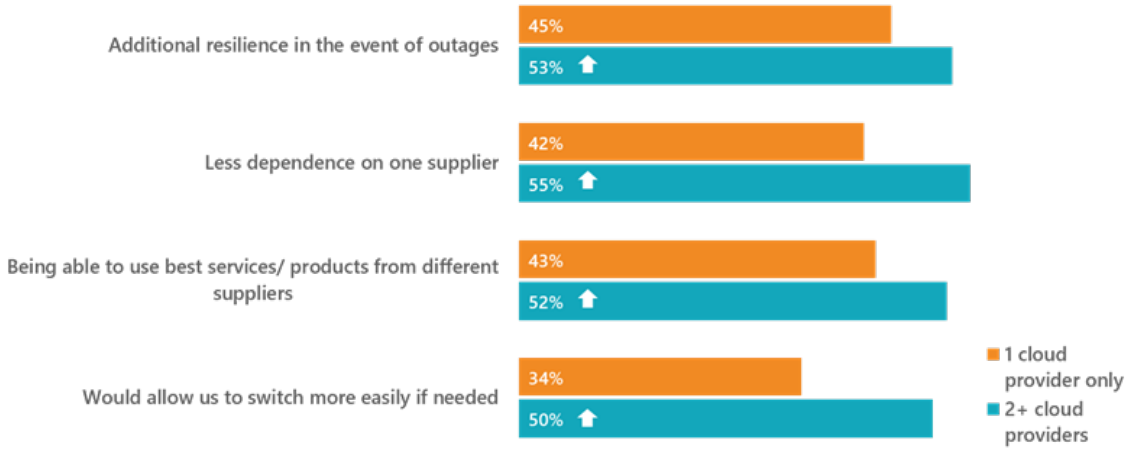

Source: Context Consulting market research report, slide 78.

[184] Proportion of all respondents.
[185] Context Consulting market research data tables, Q32.

A3.31 In the market research, 46% of those who used two IaaS/PaaS providers, and 52% of those who used three or more IaaS/PaaS providers cited "using a multi-cloud strategy where workloads of an individual app are run in more than one cloud and integrated" as an action taken to mitigate the potential for cloud lock-in.[186]

A3.32 For comparison, respondents in the Foundry 2022 Study reported the following potential benefits to a multi-cloud architecture: the agility-enhancing benefits of avoiding vendor lock-in (50%), improving disaster recovery/business continuity (47%), and greater platform and service flexibility (44%). However, enterprises (in this report this means organisations with more than 1,000 employees) are more interested in platform and service flexibility (55%), while for small and medium businesses the top desire is cost savings/optimisation (52%).[187] These results are broadly consistent with what we saw in the market research.

A3.33 In the 2023 Oracle report, data sovereignty (i.e. the idea that the data organisations use is subject to the legal and regulatory regimes of the localities where it is collected, stored and analysed) is the most frequently cited motivation for multi-cloud, followed closely by cost optimisation and, more distantly, by business agility and innovation.[188] The report also categorises the drivers for multi-cloud adoption into two broad groupings:

a) **Defensive** — i.e. intended to guard against pricing or total-cost-of-ownership shifts, vendor lock-in, regulatory compliance and business continuity interruptions.
b) **Offensive** — i.e. to enable best-of-breed cherry-picking of services across cloud providers, leverage incentives and functionality offered by existing strategic vendors, and empower internal stakeholders to use their clouds of choice.

[186] Context Consulting market research data tables, Q64.
[187] https://resources.foundryco.com/download/cloud-computing-executive-summary, page 3 [accessed 22 August 2023].
[188] Multicloud in the Mainstream: Making IT Work 'As Advertised'. Commissioned by Oracle. S&P Global Market Intelligence. February 2023. Figure 1, pages 4-5.

Figure A3.3. Benefits of using multi-cloud highlighted by the 2023 Oracle report

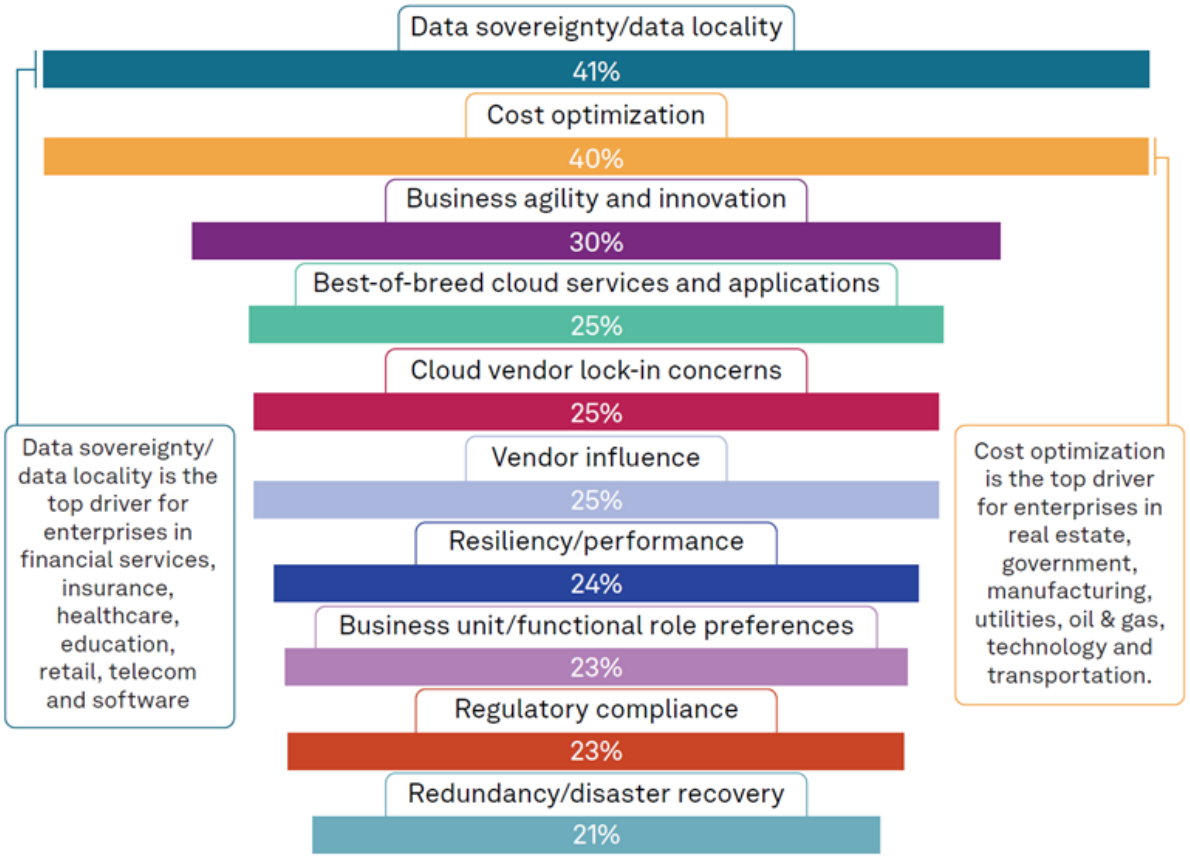

Source: Multicloud in the Mainstream: Making IT Work 'As Advertised'. Commissioned by Oracle. S&P Global Market Intelligence. February 2023.

A3.34 The market research also asked about perceived **challenges of using multi-cloud**. As shown in the chart below, the most frequently mentioned challenge is moving data across IaaS/PaaS providers.

Figure A3.4: Challenges of using multi-cloud as reported in the Context Consulting market research

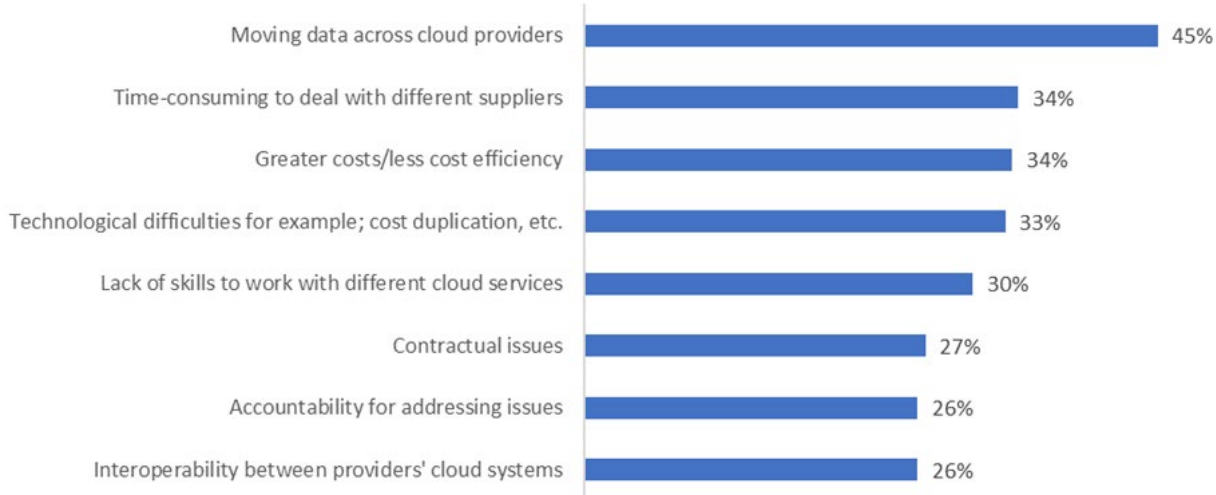

Source: Ofcom analysis of Context Consulting market research data tables, question 31.

A3.35 Comparing responses of those who already use more than one IaaS/PaaS provider with those who only use one shows a contrast between perceived versus experienced challenges.

The most frequently mentioned challenge for both groups is moving data across IaaS/PaaS providers, and the percentages are very close. For all other types of challenges, the respondents who use more than one IaaS/PaaS provider cite challenges more often than those who only use one IaaS/PaaS provider. This may suggest that some challenges are often discovered in practice and are not anticipated in advance. Figure A3.5 shows comparatively larger differences between those who use more than one IaaS/PaaS provider and those who only use one IaaS/PaaS provider for some challenges, including accountability for addressing issues (31% compared to 21% respectively), technological difficulties (37% compared to 29%), and interoperability (31% compared to 23%).

Figure A3.5: Challenges of using multi-cloud as reported in the Context Consulting market research

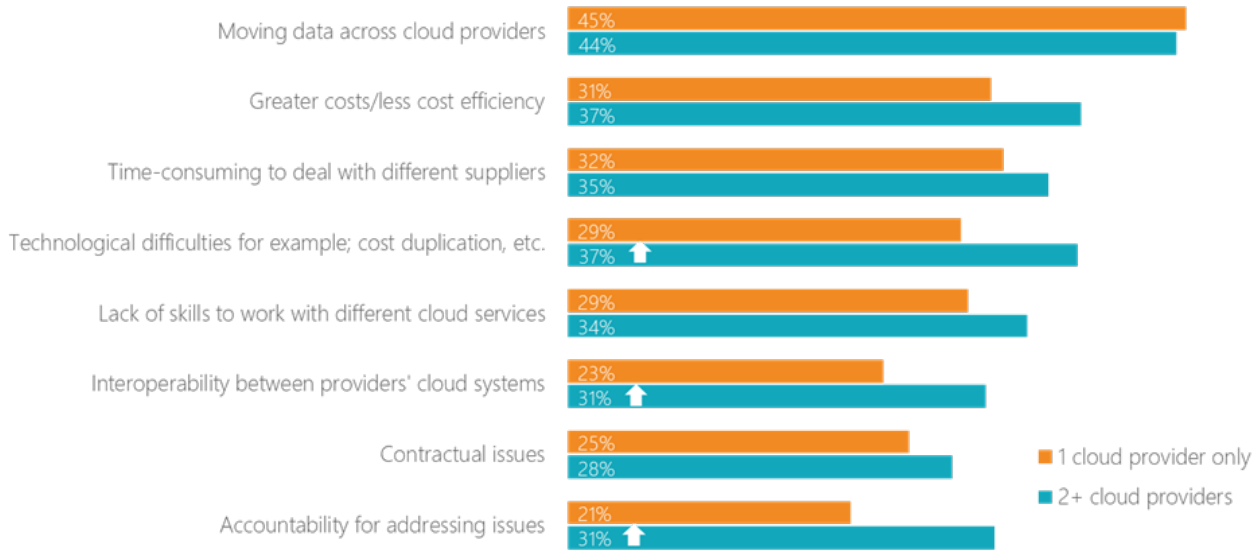

Source: Context Consulting market research report, slide 79.

A3.36 There is some variation by other respondents' characteristics, with no major difference by number of employees and some variation by industry (see Context Consulting research data tables, Q31). Some challenges are cited more often by respondents with longer experience of cloud services. Laggards (perhaps due to more limited experience) cite challenges less frequently (except for moving data).

A3.37 According to the 2023 Oracle report, the most significant challenges of multi-cloud include cloud provider management and networking/interconnectivity. The ability to manage workloads and data across multiple public cloud provider platforms emerges most often in the top-three multi-cloud challenges for enterprises (i.e. organisations with more than 1,000 employees). The report suggests that this issue essentially comes down to shortage of talent and expertise.[189]

[189] Multicloud in the Mainstream: Making IT Work 'As Advertised'. Commissioned by Oracle. S&P Global Market Intelligence. February 2023. Figure 5, p.7.

Figure A3.6: Challenges of multi-cloud according to the 2023 Oracle report

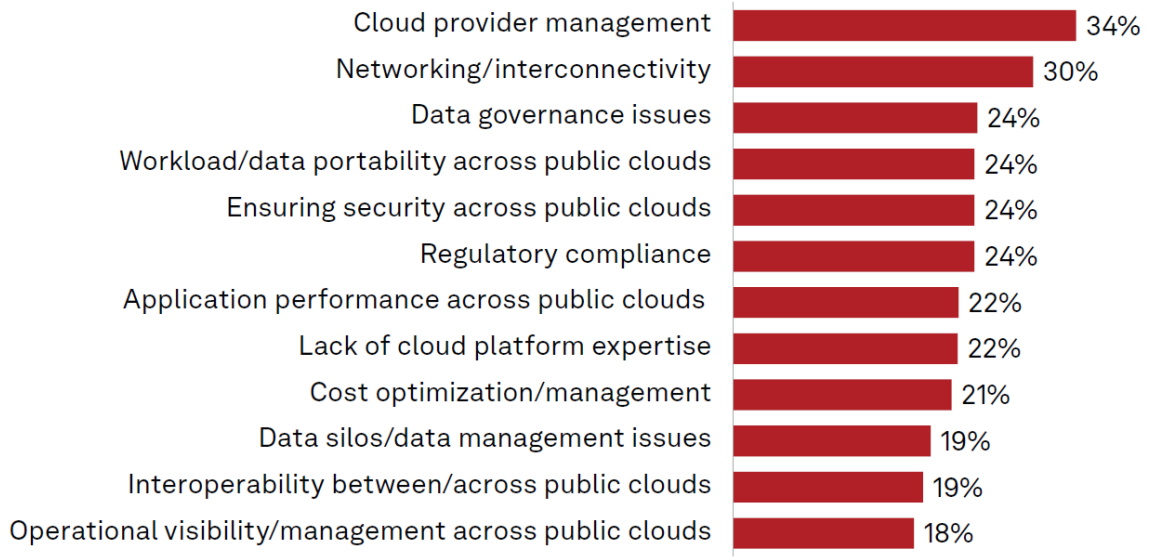

Source: Multicloud in the Mainstream: Making IT Work 'As Advertised'. Commissioned by Oracle. S&P Global Market Intelligence. February 2023.

A3.38 We also asked some customers to identify obstacles to the adoption of multi-cloud strategies. One customer ([✂]) told us that it would like to have the choice to operate a multi-cloud strategy as this would give it options in terms of innovation and enable it to adopt the best technology solutions for its needs. It said it would also provide a better commercial leverage, to respond to price increases but also in relation to other commercial relationships with cloud providers.[190] The primary factor limiting this customer's ability to adopt multi-cloud architecture for the same workloads was the cost to train and hire staff. The customer also noted limitations to, or a lack of, interoperability was important. In addition, moving workflows out of their current cloud provider to another cloud provider would make it harder for it to reach spending commitments. The customer also suggested that data egress fees are another factor limiting its ability to operate multi-cloud architecture.[191] As noted above, [✂] told us that the use of integrated multi-cloud necessarily includes additional complexity in the workload design and reduces optimisation in use of the underlying cloud services.[192] Another customer, [✂], noted that, although it duplicates its databases on two cloud providers, it would be cost prohibitive to maintain both clouds actively all the time – and, therefore, the back-up cloud of one provider is more basic.[193]

A3.39 In general, cost, increased complexity, need for skilled staff and data egress fees were mentioned frequently in the customer responses as the obstacles or challenges to multi-cloud use.

A3.40 The Public First's survey also asked IaaS/PaaS users (including those who only use one cloud provider) how easy or difficult they would find it to add and integrate and an additional cloud infrastructure provider to their current setup. The majority (62%, i.e. 444 out of 716 respondents) said that they would find it easy, compared to 21% (i.e. 150 out of 716

[190] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[191] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[192] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[193] [✂].

respondents) who said that it would be neither easy nor difficult and to 15% (i.e. 107 out of 716 respondents) who said that it would be difficult.[194]

A3.41 We note that customers responding to the Public First survey may not have answered this question specifically with integrated multi-cloud in mind (as we have defined it in Section 4), and may also (and perhaps more often) have referred to the ease of adding cloud infrastructure providers requiring more limited integration efforts, such as when a customer is adding an ISV to its existing architecture on the public cloud it is already using; adding a second cloud provider for limited and/or specialised workloads which require limited integration (i.e. at the more siloed end of the multi-cloud spectrum); and adding a public cloud provider to its on-premises solution or private cloud. In addition, the question was also posed to customers who are currently using only one cloud infrastructure provider. As these customers do not have experience of multi-clouding, they may have underestimated the challenges of adding and integrating an additional public cloud provider.

Evidence from [✂] on multi-cloud – and our assessment

A3.42 [✂] submitted a quantitative analysis of its UK customers' usage and opportunity data. According to [✂], the results of this analysis show that "[✂]% of customers (weighted by revenue) that appear in [✂] awarded tenders to at least one other cloud service provider between 2017-22, implying that they were multi-clouding."[195]

A3.43 Moreover, as additional evidence of the prevalence of multi-clouding among its UK customers,[196] [✂] stated that the distribution of the revenue share of UK customers in the [✂] by the number of cloud service providers they awarded tenders[197] to between 2017 and 2022 is the one depicted in the figure below. The figure shows that more than [✂]% of [✂] UK customers (in terms of revenues) had awarded tenders to one cloud provider only (i.e. [✂]), [✂]% had awarded tenders to one other cloud provider (apart from [✂]), [✂]% had awarded tenders to two other cloud providers (apart from [✂]) and the remaining [✂]% had awarded tenders to three or more other cloud providers (apart from [✂]).[198]

Figure A3.7: Distribution of the revenue share of customers in [✂] by the number of cloud service providers they awarded tenders to between 2017 and 2022

[✂]

Source: [✂].

[194] See question 55 of the Public First's survey [accessed 7 September 2023].
[195] [✂]. According to [✂], the [✂] quoted in the text might underestimate the total proportion of its customers who use multi-cloud or hybrid solutions, because many customers may have acquired IT services from multiple cloud service providers using processes other than tendering; [✂] may not have participated in all tenders issued by customers; in some case, the identity of the cloud service provider to which the tender was awarded is not known; and some customers may have bought cloud services outside the period included in the analysis.
[196] [✂].
[197] We note that, while [✂] uses the term "tender" in its submission, its dataset lists commercial opportunities, which may not always be formal tenders or negotiation processes.
[198] In total, based on the figures reported in the text, almost [✂] of [✂] UK customers (in terms of revenues) had awarded tenders to at least one other cloud provider (apart from [✂]), during the period under consideration. The reason for the discrepancy with the [✂]% reported in paragraph A3.42 appears to be the different ways in which [✂] counts multi-clouding revenues. In particular, [✂].

A3.44 For the reasons set out below, we have concluded that [✂] arguments do not materially alter our assessment of the prevalence of multi-cloud.

A3.45 First, [✂] approach – based on weighing customers by their spend on [✂] – does not provide a meaningful measure of the prevalence or extent of multi-clouding. This is because it only takes into account how much [✂] customers spent on [✂] services during the period under consideration – without considering how much [✂] customers may have spent on other cloud providers (which [✂] may not know). A simple example illustrates this concern with [✂] approach: if [✂] had only one customer (which would, as a result, account for 100% of [✂] revenues) and this customer were to use another cloud provider for a new workload (irrespective of whether this workload were to be small or large), [✂] indicator of the prevalence of multi-cloud would still be equal to 100%. We think the indicator proposed by [✂] fails to reflect how much [✂] existing customers are spending with rival cloud providers and is therefore not a good measure of how extensive multi-cloud may be. In practice, our qualitative evidence[199] suggests that [✂] large customers typically only use secondary providers for a minority of workloads. This would suggest that the fraction of [✂] customers' revenue that it has lost to rival cloud providers could be limited.

A3.46 Second, [✂] approach includes all sales of a customer who it records as having allocated an opportunity to a rival cloud service provider. It therefore implicitly assumes that all of these workloads face competition from rival providers. However, our engagement with customers suggests that customers who multi-cloud tend to do this only for workloads which are sufficiently siloed from the architecture on their primary cloud provider.[200] As such, assuming that all workloads of a multi-clouding customer are "up for grabs" is likely to overstate the constraint of multi-clouding.

A3.47 Third, by using customer revenues as weights, the indicator used by [✂] gives more weight to large customers, i.e. those who spent the most on [✂] cloud services. This is likely to overstate the actual prevalence (as measured by the fraction of [✂] customers that multi-cloud). To assess the robustness of [✂] approach in this respect, we have calculated the unweighted share of [✂] customers who also allocated workloads to other cloud providers during the period under consideration. This share is equal to [✂]% of [✂] customers,[201] well below the [✂]% indicated by [✂] and also below what our market research suggests (see above in this annex).[202]

A3.48 Fourth, we are concerned that the datasets used by [✂] to calculate its indicator of the prevalence of multi-cloud present several flaws and cannot be relied upon. According to [✂], the main dataset used in its quantitative analysis (the [✂] dataset) consists of data recorded in [✂], a customer relationship management ('CRM') platform, where "opportunities inputs are manually made by members of [✂] sales teams. As a result, the

[199] See also Section 5 and Section 6.
[200] See paragraphs A3.6 to A3.15 above.
[201] Ofcom's analysis based on data used for [✂]. There are [✂] customers in the dataset, of which [✂] multi-clouding (which corresponds to [✂]% of the total).
[202] Ofcom's analysis based on data used for [✂]. As part of our robustness checks of [✂] results, we have also calculated the unweighted distribution of customers in [✂] by the number of cloud providers they awarded tenders to between 2017 and 2022, to compare it with the distribution of the revenue share calculated by [✂] and shown in Figure A3.7 above. Our calculations show a more limited prevalence of multi-cloud than in [✂] analysis, e.g. [✂]% of [✂] customers had awarded tenders to one cloud provider only (i.e. [✂]), [✂]% had awarded tenders to one other cloud provider (apart from [✂]) and the remaining [✂]% had awarded tenders to two or more other cloud providers (apart from [✂]).

data is not always comprehensive and may contain errors and inconsistencies."[203] [✂] also noted that the data compiled by members of its sales teams "is based on manual input and is often comprised of anecdotal feedback from the customer, particularly in relation to competitor information."[204] Lastly, [✂] noted that its [✂] dataset "does not cover all [✂] customers. [✂] is therefore not able to accurately assess how representative and comprehensive the [✂] dataset in percentage terms."[205]

A3.49 Overall, we consider that those arguments do not materially alter our assessment about the level of multi-clouding in cloud infrastructure services.

A3.50 In a separate submission, [✂] submitted a quantitative analysis based on its customers' committed spend and opportunity data. According to [✂], over [✂]% of [✂] customers with committed spend contracts in the UK active between [✂] awarded opportunities to at least one other cloud provider during or before the period they had a contract with [✂]. Moreover, according to [✂] calculations, committed spend from these multi-clouding customers would account for over [✂]% of [✂] total committed spend from contracts active in this period. According to [✂], these results are consistent with evidence of multi-cloud, showing no evidence of customers with committed spend discounts being incentivised to 'single source'.[206]

A3.51 In respect of this analysis, we note that [✂] customers with committed spend contracts are relatively large and need to commit a significant amount of minimum spend to qualify for discounts.

A3.52 In particular, as [✂] told us, for "a very small minority of customers who have business objectives that require [✂] services over time and who can broadly predict their minimum service needs across services, [✂] offers additional discounts in exchange for a commitment to spend a certain amount across most [✂] services."[207] Specifically, customers "can choose to make a spend commitment that is appropriate for their needs, and can commit for one year or for multiple years as they wish (subject to a minimum annual commitment of [✂]), or they can decide not to make a commitment and instead use [✂] services under public pricing."[208] Moreover, in addition to customers entitled to discounts which "depend on the commitment made by the customer and [✂]",[209] the sample used by [✂] includes [✂].

A3.53 We agree with [✂] that large customers who have committed spend discounts can in some cases use more than one cloud provider – this is consistent with our own outreach to large customers. However, this multi-clouding often involves the use of a primary provider for a large majority of their cloud workloads. These customers only use secondary providers for a minority of workloads, which are sufficiently siloed from their main architecture on their primary provider.[210]

A3.54 In addition, the data used by [✂] in its quantitative analysis (notably, the [✂] dataset) is the same as the data used for the multi-cloud analysis of all its UK customers, as discussed

---

[203] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[204] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[205] [✂] response dated [✂] to the s.174 notice dated [✂], [✂].
[206] [✂].
[207] [✂].
[208] [✂].
[209] [✂].
[210] Ofcom / [✂] meeting, [✂]. Ofcom / [✂] meeting, [✂]. [✂] response dated [✂] to our customer questionnaire.

above. As such, the same reservations apply about the quality and accuracy of the information included in this dataset expressed in paragraph A3.48 above – which [✂] also acknowledges.

A3.55 Overall, we conclude that the analysis of [✂] for committed spend discount customers does not disprove our concern that committed spend discounts can encourage such customers from placing most or all of their cloud needs with [✂].

Evidence on switching

Prevalence of switching

A3.56 The market research asked IaaS/PaaS users (i.e. 'considerers' were excluded for this question) whether they have switched IaaS/PaaS providers in the past completely (and stopped using the previous IaaS/PaaS provider), whether they have taken on an additional IaaS/PaaS provider in the past, and whether they considered switching but did not switch. Overall, **18% of respondents said they switched IaaS/PaaS providers in the past, 35% said they took on an additional IaaS/PaaS provider**, 35% considered switching but did not switch, and 23% never considered switching. If we exclude respondents using only private cloud, 21% said they switched IaaS/PaaS providers in the past, 39% said they took on an additional IaaS/PaaS provider, 32% considered switching but did not switch, and 21% never considered switching.

A3.57 There are differences in the proportions of respondents who switched according to their attitude to technology adoption (see Figure A3.8). Early adopters of technology were more likely to have switched (34%) and to have taken on new IaaS/PaaS providers (47%), while laggards were less likely to have done either of these (12% and 20%, respectively), and less likely to have considered switching (49% had never considered).

Figure A3.8: Switching behaviour in the Context Consulting market research by attitude to technology

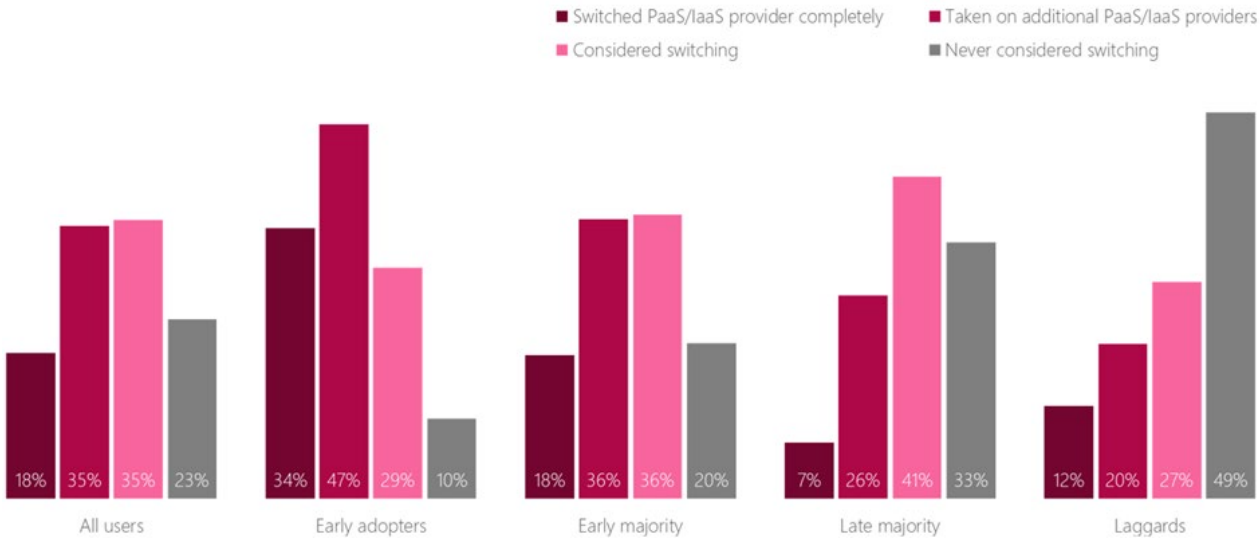

Source: Context Consulting research report, slide 105.

A3.58 The following groups were also **more likely to have switched** in the past compared to an average respondent (18%): companies in IT & Technology (25%), organisations younger than

2 years (25%), AWS users (26%) and Google users (23%), and organisations that used 3 or more IaaS/PaaS providers (36%). The higher proportion of switchers in these categories can be explained by their attitude to technology, similar to "early adopters", or their specific needs.

A3.59 At the same time, the following groups were **less likely to have switched** (than the average respondent at 18%): smaller organisations with 10-49 employees (at 13%), public sector (9%), organisations older than 20 years (9%), that did not use hyperscalers (13%), and organisations that used only one IaaS/PaaS provider (11%). Public sector respondents were more likely to have never considered switching (at 31%). There is an overlap of some customers in these categories with the 'laggards'[211] category in terms of technology adoption which could mean they have yet to reach the end of their contract, or may not be very 'tech savvy', or may not have sufficient internal resources to invest in searching and investigating the market and evaluating options.

A3.60 It is difficult to find publicly available information from the UK or other countries to put our findings into context and compare them. The 2022 Report of the Japanese Fair Trade Commission cites results of their survey, which similar to the Context Consulting market research was aimed at IaaS and PaaS, and which finds that 15.7% of the respondents (86 companies among 548) switched cloud provider in a previous decade.[212] This number is similar to our finding of 18% of respondents who switched a IaaS/PaaS provider.

A3.61 The qualitative part of the Context Consulting market research,[213] which captured a richer and more in-depth understanding of customer responses, suggests that the level of switching found in the quantitative part of the market research could be over-stated, perhaps due to misunderstanding of the question, or lack of ability of respondents to explain their more nuanced experience:

a) We encountered few if any examples of organisations switching away from one of the hyperscalers.
b) The switching we found was typically from a data centre to an IaaS/PaaS environment.
c) In some cases, firms were adding additional platforms (e.g. bringing AWS onboard alongside Azure).
d) It was still relatively early in the adoption journey for most companies, and they were evaluating progress rather than looking to make significant changes.
e) In most cases, firms were still on the way in, not out, of their IaaS/PaaS environments.[214]

A3.62 Overall, it is possible that the respondents to the quantitative part of the market research were conflating a migration from on-premises IT/private cloud to public cloud, switching between first- and third-party services within a public cloud and switching between public clouds.

A3.63 The qualitative part of the market research also suggests that there is a lot of inertia reducing incentives to switch and companies are eager to avoid the disruption this would

[211] The market research asked respondents to indicate how they would best describe their business attitude towards investing in technology. In response, some classified their business as ones that "only buys established and proven technology that is the standard for our industry" OR one that "waits as long as possible before investing in new technology". Laggards refers to these two categories.
[212] Report on Trade Practices in Cloud Services Sector. June 2022, page 49. Available at https://www.jftc.go.jp/en/pressreleases/yearly-2022/June/221102EN.pdf [accessed on 22 August 2023].
[213] Context Consulting market research report, slide 110.
[214] Context Consulting market research report, slide 110.

entail. Many decision-makers acknowledged that a de facto lock-in exists and this was often perceived to be a function of internal factors rather than provider-imposed restrictions. Given the inertia, something very significant would need to happen to prompt a switch away from a cloud provider, such as a substantial price hike, a deterioration in technical performance, or significant security concerns.

A3.64 Responses to our customer questionnaire show that a "complete" switch of cloud providers is very rare. Out of all respondents, only [✂] told us they switched providers (from [✂] to [✂] and [✂]), and the switch took place for a number of reasons, including technical capability and scalability, speed of innovation and costs.[215] Other respondents only moved a few workloads from one cloud provider to another (i.e. a partial switch), which indicates that customers are avoiding a full switch and are only looking to switch where necessary (e.g. to concentrate workstreams with a single cloud provider to avoid egress fees). For example, [✂].[216]

A3.65 Some customers (e.g. [✂]) told us they simply have not had a compelling reason to switch so far, especially given the cost, time and effort that such a move would entail.[217]

A3.66 In the survey conducted by Public First,[218] IaaS/PaaS users were asked whether they had ever switched cloud infrastructure provider in the past. The results indicate that switching is limited, since only 26% of IaaS/PaaS users (i.e. 185 out of 716 respondents) said that they had switched provider. In contrast, 35% of IaaS/PaaS users (i.e. 251 out of 716 respondents) had considered switching, but had not switched in the end; and 36% of IaaS/PaaS users (i.e. 258 out of 716 respondents) had not considered switching.

A3.67 Public First also asked all 716 IaaS/PaaS users whether they thought they would be likely to change their main cloud infrastructure providers to a different IT services provider in the next few years.[219] According to the survey results, 51% of IaaS/PaaS users (i.e. 364 out of 716 respondents) reported that they were likely to switch cloud infrastructure provider in the future, compared to 20% of IaaS/PaaS users (i.e. 143 out of 716 respondents) who were unlikely to switch in the future. A further 29% of IaaS/PaaS users (i.e. 208 out of 716 respondents) said that they were neither likely or unlikely to switch in the future. We note that these answers may be inaccurate, since IaaS/PaaS users were asked about their intentions to switch in the future.[220]

A3.68 Furthermore, Public First asked those 185 IaaS/PaaS users who had switched their cloud infrastructure provider in the past whether they had switched to another cloud

---

[215] [✂] response dated [✂] to our customer questionnaire, questions [✂].
[216] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[217] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[218] See question 29 of the Public First's survey [accessed 7 September 2023].
[219] See question 35 of the Public First's survey [accessed 7 September 2023].
[220] Public First also asked the 143 IaaS/PaaS users who said they were unlikely to switch cloud provider in the future about the possible reasons for this. 80% of this group (i.e. 114 out of 143 respondents) said that they were satisfied with their current cloud infrastructure provider; 11% (i.e. 16 out of 143 respondents) said that it would be too complicated to change their main provider; and 8% (i.e. 11 out of 143 respondents) said that it would be too expensive. See question 36 of the Public First's survey [accessed 7 September 2023]. We note that these results may not entirely be representative, due to the small number of respondents, compared to the overall size of the sample (1,001 respondents).

infrastructure provider or to an on-premises solution. The vast majority – 69%, i.e. 127 out of 185 'switchers' – had switched to another cloud infrastructure services provider, whereas 29% – i.e. 54 out of 185 "switchers" – had moved to an on-premises solution.[221]

A3.69 We note that it is not entirely clear what type of switching these questions may capture. Given that 32% of IaaS/PaaS respondents to Public First's survey reported using private cloud only and 48% reported using public and private cloud (as noted in paragraph A3.13 above), the switching that IaaS/PaaS users reported may have been between private and on-premises solutions. In addition, the questions may capture switching within public clouds (i.e. between first and third parties hosted on the same cloud) rather than between clouds. Moreover, the question may capture partial switching (i.e. of a few workloads only), as opposed to changing IaaS/PaaS provider for all workloads.

Reasons for switching and switching experience

A3.70 The market research asked those respondents who switched (164 respondents) an open question on what the **reason for changing the IaaS/PaaS providers** was. Looking at verbatim responses, we can see that more than half of the responses mentioned price, cost or "value for money". About a third mentioned customer service, quality of service or various technological aspects.

A3.71 We also asked those who switched IaaS/PaaS provider whether they found the experience **difficult or easy**. About a half (47%) of switchers described the process as very easy or quite easy, a quarter found it neutral, and 29% found it very difficult or quite difficult.[222]

Figure A3.9: Switching experience in the Context Consulting market research

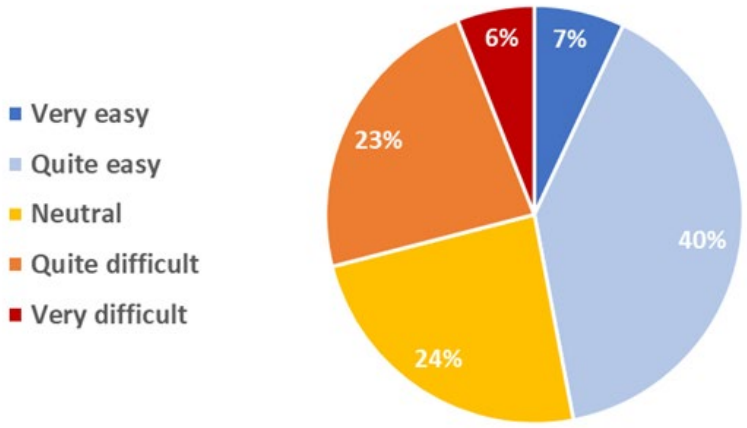

Source: Context Consulting market research report, slide 109.

A3.72 We asked those who **added an IaaS/PaaS provider** (308 respondents) about their experience. 12% said it was very easy, 43% - quite easy, 25% - neutral, 17% - quite difficult

[221] See question 31 of the Public First's survey [accessed 7 September 2023].

[222] These results are similar to those found by Public First in their survey [accessed 7 September 2023], question 28. Asked how easy or difficult it would be to switch cloud infrastructure provider, 54% of the 1,001 respondents said that it would be easy; 27% said that it would be difficult; and 17% said that it would be neither easy nor difficult. Note that this question asks about a hypothetical switching and was answered by all participants to the survey, including those who had not switched cloud infrastructure provider in the past.

and 3% - very difficult. Overall, the process of adding an IaaS/PaaS provider seems to have been less challenging than switching an IaaS/PaaS provider completely.[223]

A3.73 We also asked respondents what would need to happen to prompt them to switch in the future. The top three reasons are better service quality (at 48%), lower price (46%) and improved level of security (41%). Answers to this hypothetical question put more emphasis on quality and security as opposed to price, compared to the question about the reasons for an actual switch in the past where price was the most frequently cited reason.

Figure A3.10: Potential reasons for switching providers

Source: Context Consulting market research report, slide 110.

A3.74 We also asked what would prompt cloud users to **add an additional IaaS/PaaS provider**. Cost factors (better price and better value for money) were most important, followed by service quality and security.

[223] Context Consulting market research data tables, Q50.

Figure A3.11: Potential reasons for adding an IaaS/PaaS provider.

Source: Context Consulting market research report, slide 111.

Challenges and obstacles to switching

A3.75 The most frequently cited challenges of switching in the market research were time and cost (43% overall), the need to retrain staff (33%), followed by technical difficulties in transferring applications and software and data. However, when asked about the main challenge, the need to retrain staff drops to fourth place after time and cost and technical difficulties.

Figure A3.12: Perceived challenges of switching in the Context Consulting market research

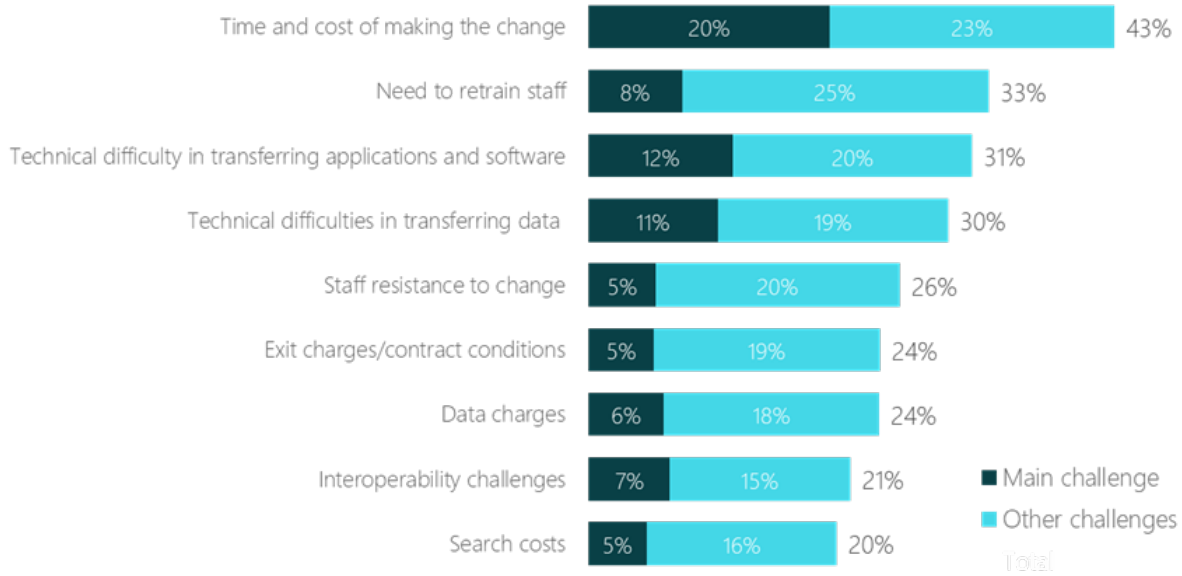

Source: Context Consulting market research report, slide 123.

A3.76 There are some differences in the order of importance of the challenges depending on whether the respondents have switched before, considered switching or not. While "time and cost" was the top challenge across all such categories of respondents, those who have

switched or taken on an additional IaaS/PaaS provider viewed technical issues as the second most important challenge. Those who considered switching but did not, and those who never considered switching, viewed the need to retrain staff as the second most important challenge. The need to retrain staff was the third most important challenge for those who switched in the past and those who have taken an additional IaaS/PaaS provider.

A3.77 The market research also asked what strategy, if any, customers have taken to mitigate risks of cloud lock-in. Most respondents said they have taken some action to mitigate the potential for cloud lock-in, with "ensuring data portability" being the most common. However, as discussed above and further detailed in Section 5, the evidence we have received indicates that – despite use of these mitigation strategies – obstacles to switching are likely to remain material for some customers and use-cases.

A3.78 Customers who engaged with us also often mentioned the time and cost involved in migrating applications and data, with some specifically emphasising egress fees ([✂] and [✂]), and need to retrain staff.[224] For example, [✂] explained that it started using AWS's public cloud services for its digital service (website, app, backend etc.) in 2016/17. It considers that it could not switch to Google or Azure at present because it is "a bit tied" to AWS for technical reasons and due to the cost of retraining all its developers.[225]

A3.79 In the survey conducted by Public First,[226] IaaS/PaaS users were asked whether their cloud infrastructure provider offers dedicated features to support switching away from their services. The majority of IaaS/PaaS users (51%, i.e. 365 out of 716 respondents) responded in the affirmative, whereas the remaining 24% of IaaS/PaaS users (i.e. 172 respondents) said that their cloud infrastructure provider did not offer support to switch away and 25% of IaaS/PaaS users (i.e. 179 respondents) did not know.

A3.80 As noted in paragraph A3.66 above, the results of the survey conducted by Public First indicated that only a minority of IaaS/PaaS users (namely, 185 out of 716 respondents) had ever switched cloud infrastructure provider in the past.[227] Within this group of "switchers", the vast majority – 79%, i.e. 146 out of 185 respondents – had made use of the dedicated features offered by their cloud infrastructure provider to make switching easier.[228]

A3.81 Asked about the relative proportion of financial costs for switching, 54% of the "switchers" (i.e. 100 out of 185 respondents) said that the majority or vast majority of costs were internal (e.g. retraining of staff and creation of new systems). In contrast, 15% of the "switchers" (i.e. 28 out of 185 respondents) said that the majority or vast majority of costs were external (e.g. data transfer fees); 25% of the "switchers" (i.e. 46 out of 185

[224] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our cloud services questionnaire, question [✂].
[225] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].
[226] See question 30 of the Public First's survey [accessed 7 September 2023].
[227] Question 29 of the Public First's survey [accessed 7 September 2023].
[228] Question 32 of the Public First's survey [accessed 7 September 2023]. We note that these results may not entirely be representative, due to the small number of respondents, compared to the overall size of the sample (1,001 respondents). The small number of switchers (35 respondents) who had not made use of the dedicated features offered by their cloud infrastructure provider to make switching easier provided different reasons for their choice, such as: Those features not helpful for the respondent's business; features not offered by the cloud provider; features being too expensive; and lack of awareness about the features. See question 33 of the Public First's survey [accessed 7 September 2023]. We note that these results may not entirely be representative, due to the small number of respondents, compared to the overall size of the sample (1,001 respondents).

respondents) said that costs were around equal between internal and external; and 5% of the "switchers" (i.e. 9 out of 185 respondents) said that there were no significant financial costs associated with switching.[229]

A3.82 Public First asked the same question – hypothetically – to all 716 IaaS/PaaS users; the results were slightly different. Specifically, 46% of IaaS/PaaS users (i.e. 329 out of 716 respondents) said that the majority or vast majority of costs were internal (e.g. retraining of staff and creation of new systems). In contrast, 15% of IaaS/PaaS users (i.e. 107 out of 716 respondents) said that the majority or vast majority of costs were external (e.g. data transfer fees); 28% of the IaaS/Pass users (i.e. 200 out of 716 respondents) said that costs were around equal between internal and external; and 6% of the IaaS/PaaS users (i.e. 43 out of 716 respondents) said that they did not believe that there would be significant financial costs for switching.[230] We note that these answers may be inaccurate, since only a minority of the 716 IaaS/PaaS users have had any actual experience of switching.

Evidence from [✂] on switching – and our assessment

A3.83 In response to our interim report,[231] [✂] submitted a quantitative analysis of its customer data, which, according to [✂], would indicate that [✂] customers can and do switch.

A3.84 Specifically, [✂] told us that, in each of [✂] and [✂], approximately [✂] of its UK customers "churned" [232] from [✂] on an annual basis, considering all cloud services offered by [✂]. This share would be approximately equal to [✂] (again, in each of [✂] and [✂]), if one only considers cloud compute services in the analysis.[233] According to [✂], these figures would exceed the churn rate from traditional on-premises IT providers.

A3.85 In addition, [234] [✂] compared the level of spend of its UK customers between [✂] and [✂].[235] According to [✂] analysis, across all cloud services, [✂]% of [✂] customers (accounting for [✂]% of [✂] revenues) decreased their spend between the two periods. This proportion is equal to [✂]% (accounting for [✂]% of [✂] revenues), if only cloud compute services are included in the analysis. According to [✂], these figures indicate that customers can move significant workloads (and so reduce their spend on [✂] services in the process) from [✂] to other cloud providers, which would be evidence of significant switching.

---

[229] Question 34 of the Public First's survey [accessed 7 September 2023]. We note that these results may not entirely be representative, due to the small number of respondents, compared to the overall size of the sample (1,001 respondents).
[230] Question 37 of the Public First's survey [accessed 7 September 2023].
[231] [✂].
[232] In [✂] analysis, based on data about UK revenues and usage, a customer is churning if the following two conditions are both met: i) the customer reduces their spend on [✂] services by at least 80% for three consecutive months; and ii) the spend does not return to 80% of the original level within six months. As [✂] also notes, although these two conditions together imply a significant reduction in spend which may constitute switching, they do not necessarily guarantee that the customer has switched away from [✂]. In addition, the sample used by [✂] includes customers with positive data transfer-out spend or usage.
[233] In [✂] analysis, the [✂]% quoted in the text above would correspond to over [✂] customers across all its cloud services, whereas the [✂]% would correspond to over [✂] customers using its cloud compute services.
[234] [✂].
[235] The sample used by [✂] includes customers with positive data transfer-out spend or usage and who also spent at least [✂] US dollar on [✂] services in [✂]. In addition, [✂] excludes the top [✂]% of customers from their analysis because "[✂]% of these customers multi-cloud based on [✂] and are more likely to switch workloads rather than churning outright." See [✂].

A3.86 To start with, we note that the churn rates calculated by [✂] ([✂]% across all services and [✂]% for cloud compute services only) are similar to those found by the Context Consulting market research and consistent with a finding that switching levels are low in the cloud infrastructure services market.[236]

A3.87 Furthermore, we note that, as far as churn rates are concerned, the approach used by [✂] to identify "churning" customers – defined as customers reducing their spend by at least 80% for three consecutive months and not returning to 80% of their original spend within six months (see footnote 232) – may only imperfectly capture the actual churning levels in the cloud market. This is because, for example, a customer may still continue to buy from [✂] and/or significantly increase its spending on [✂] services after six months. [✂] acknowledges that its measure is imperfect.

A3.88 In addition, we have concerns about the robustness of [✂] results, when the data are analysed at a more granular level. The analysis carried out by [✂] includes many small customers who only spend a minimal amount on [✂] cloud services during the entire period considered (and so, in our view, may have just tried [✂] services or only used those services occasionally). Once these customers are removed from the sample, [✂] churn rates become much smaller. By way of example, removing customers who spent less than $[✂] during the entire period under consideration makes the churn rate calculated by [✂] drop from [✂]% to [✂]% in [✂] and to [✂]% in [✂], across all services. Removing customers who spent less than $[✂] during the entire period would have an even bigger effect, since the churn rates would be equal to [✂]% in [✂] and to [✂]% in [✂] (again, across all services).[237]

A3.89 We also found that customer spend in [✂] data is heavily skewed towards the largest customers. For example, the top [✂]% of [✂] customer accounts for [✂]% of [✂] total revenues in the dataset – and the top [✂]% for [✂]% of the total revenues.[238] This is also reflected in the churn rates calculated for different customer spend deciles. For example, for customers who are above the [✂] decile in terms of spend, the churn rate is below [✂]% in each decile, in both [✂] and [✂]. In fact, the top [✂]% customers in terms of spend have churn rates of approximately [✂]%, i.e. they have practically never switched (using [✂] definition of switching) during the period under consideration.[239]

A3.90 As far as the comparison of spend levels between [✂] and [✂] carried out by [✂], we have also found those to lack robustness. First, we note that a reduction in spend alone is not a meaningful measure of switching. Second, when we include in the analysis the top [✂]% of [✂] customers (which [✂] had not included, see footnote 235), the share of [✂] customers who decreased their spending from one period to the other is [✂]% (compared to the [✂]% calculated by [✂]), but these customers account for a very small share of total revenues, i.e. [✂]% (compared to the [✂]% calculated by [✂]), across all services.[240] In our view, this implies that, although almost [✂] of [✂] customer base in the UK has decreased

[236] See paragraph A3.56 above.
[237] Ofcom's analysis based on data used for [✂]. If we only consider cloud compute services, the churn rates would be equal to [✂]% in [✂] and to [✂]% in [✂], if we remove customers spending less than $[✂]; and to [✂]% in [✂] and to [✂]% in [✂], if we remove customers spending less than $[✂].
[238] Ofcom's analysis based on data used for [✂].
[239] Ofcom's analysis based on data used for [✂].
[240] If we include the top [✂]% of [✂] customers and only consider cloud compute services, the share of customers who decreased their spend from [✂] to [✂] is equal to [✂]% of the total number of customers, accounting for [✂]% of [✂] total revenues.

their spend between [✂] and [✂], the share of revenues they account for is actually very small, being equal to [✂] of [✂] total revenues across all services.

A3.91 Overall, for the reasons listed above, we conclude that the evidence presented by [✂] does not alter our conclusions on switching levels in cloud infrastructure services in the UK.

A3.92 As noted in paragraphs A3.66 and A3.68 above, the Public First's survey results indicate that only a minority of IaaS/PaaS users – 26%, corresponding to 185 out of 716 respondents – had switched cloud infrastructure provider in the past few years.

A3.93 Within this group of "switchers", the survey found that the vast majority – 69%, i.e. 127 out of the 185 "switchers" – had switched to another cloud provider, whereas 29% – i.e. 54 out of the 185 "switchers" – had moved to an on-premises solution.

A3.94 According to [✂], these survey results demonstrate the substantial competitive pressure of on-premises solutions on cloud service providers.[241]

A3.95 We disagree with [✂] on this point. As [✂] also notes, this is historical evidence on switching and does not represent customer switching "at the margin" in response to a small change in relative price or quality. As such, we think that this evidence is not suitable to shed light on whether on-premises solutions pose a substantial competitive constraint on cloud providers. In fact, the survey results themselves indicate that respondents use cloud services and on-premises solutions for different reasons.[242] Moreover, as noted in paragraphs A3.14 and A3.69 above, it is not clear what type of switching the Public First's survey captures. Given the relatively high proportion of IaaS/PaaS users who report to be using private cloud only or a mix of private and public cloud, some of the switching captured by the survey may be between private and on-premises solutions.

A3.96 We also consider that other results of the Public First's survey are consistent with there being a structural shift towards public cloud and with a finding that on-premises solutions do not pose a substantial constraint on cloud. In particular, out of the 716 IaaS/PaaS users, only 19 (i.e. 2%) reported having decreased their spend on cloud services. This compares with 70% of respondents who reported increasing their spend on cloud services and 27% who reported having kept spend about the same (for a combined 97%).[243] This shows that there is very little movement of customers from cloud to on-premises solutions.

A3.97 Furthermore, among the 19 IaaS/PaaS users who reported having decreased their spend on cloud services, 53% said that they had done so because their company had moved back to on-premises. With 63%, however, the most frequently stated reason was an overall reduction in the need for IT services.[244] This indicates that customers who have reported switching to an on-premises solution are not switching all/most of their spend. Rather, it is likely to reflect customers who are in the process of migrating to cloud and choosing to assign some workloads to their current on-premises provider (based on functionality considerations, in addition to price consideration).

---

[241] [✂].

[242] Respondents reported the following reasons: i) To make sure we're not locked into one supplier; ii) Compliance or regulatory requirements; iii) Increased flexibility; iv) Performance optimisation; v) Cost optimisation; and, vi) Risk mitigation. All reasons were selected by respondents with similar frequencies. See question 58 of the Public First's survey.

[243] Question 14 of the Public First's survey [accessed 7 September 2023].

[244] See question 15 of the Public First's survey. Multiple responses were allowed.

A3.98 In its submission dated [✂],[245] as an additional point, [✂] attempts to calculate individual "switching rates" for different cloud providers, starting from the overall 69% of (i.e. 127) IaaS/PaaS users who said in the Public First's survey that they had switched to another cloud provider (rather than to an on-premises solution). In particular, [✂] splits the overall 69% across individual cloud providers in proportion to their market shares. [✂] then proceeds to compare the "switching rates" attributed to each individual cloud provider with that represented by the 54 IaaS/PaaS users who switched to on-premises solutions (i.e. 29%).

A3.99 [✂] carries out this exercise for two scenarios.[246] For the first scenario, [✂] makes use of information on the market shares of individual cloud providers retrieved from IDC.[247] For the second scenario, [✂] uses information on the market shares for IaaS/PaaS providers that we published in the interim report (which only consider IaaS and PaaS revenues). According to [✂], "under both scenarios" and "under various reasonable assumptions", "the collective competitive constraint of on-premises solutions is comparable, if not larger, than the competitive constraint imposed by any single cloud service provider including the largest providers."[248]

A3.100 [✂] has not provided us with any of the underlying data and calculations used to carry out this exercise. Notwithstanding that, we have several reservations about the validity of the results reported by [✂].

A3.101 To start with, as noted in paragraph A3.95 above, the historical evidence about switching used by [✂] is not suitable to shed light on whether on-premises solutions pose a competitive constraint on cloud providers.

A3.102 Moreover, [✂] appears to treat the different providers of on-premises solutions as if they were a single provider and imply that this fictional single provider is as strong a competitor, in customers' view, as each of the hyperscalers. In reality, this is not the case. There are several on-premises providers, each of different strength – considered individually – and it is not appropriate to look at them in aggregate while comparing them with switching to hyperscalers individually.

A3.103 Furthermore, splitting the 69% of (i.e.127) IaaS/PaaS users who switched to a cloud provider (rather than to an on-premises solution) across individual cloud providers – i.e. calculating individual "switching rates" – is essentially an impossible task without additional, more granular survey results (e.g. obtained by asking participants which cloud provider they had switched to) being available.

A3.104 Any attempt to carry out this exercise without more granular survey evidence is, therefore, arbitrary and its results will differ depending on the assumptions made. In this respect, [✂] has not provided any plausible explanation of why switching by IaaS/PaaS users to a cloud provider should occur in proportion to that cloud provider's market share; and why, in implementing this method, market shares should be weighed by cloud providers' revenues (which is the method apparently used by [✂]).

A3.105 In fact, [✂] approach appears to produce inconsistent results. In the first scenario, for example, the 69% of "switchers" is split among different cloud providers as follows: Azure

---

[245] [✂].
[246] [✂].
[247] For the first scenario, [✂] refers to this publication by IDC: https://www.idc.com/getdoc.jsp?containerId=IDC_P29737 [accessed 7 September 2023]. [✂].
[248] [✂].

[✂]%, AWS [✂]%, GCP [✂]%, Oracle [✂]%, IBM [✂]%, other cloud providers [✂]%, which, combined, total [✂]%, i.e. more than 69%. Similarly, in the second scenario, the split is as follows: Azure [✂]%, AWS [✂]%, GCP [✂]%, Oracle [✂]%, IBM [✂]%, other cloud providers [✂]%, which, together, total [✂]%, i.e. again more than 69%.[249] [✂] suggests that this is explained by the fact that it modelled different switching scenarios, although it has not provided a clear explanation of how it has done that and why individual switching rates that add to more than 69% are a relevant comparator for the level of switching that the Public First's survey shows towards on-premises solutions (i.e. 29%).

A3.106 Lastly, we note that the set of cloud providers in [✂] first scenario (whose market shares [✂] retrieves from IDC) appears to include not only IaaS and PaaS providers, but also SaaS providers. This approach is, however, inappropriate, as we have focused on cloud infrastructure services in this market study.

A3.107 For the reasons listed above, we do not accept that [✂] analysis of the Public First's survey demonstrates that there is substantial competitive pressure of on-premises solutions on cloud service providers.[250]

Microsoft's views on switching – and our assessment

A3.108 In response to our interim report,[251] Microsoft said that low levels of switching in the cloud infrastructure market are not necessarily evidence of weak competition, as it may also suggest that, for cloud customers, the gains from switching do not outweigh the costs of switching.

A3.109 Microsoft also pointed to the following characteristics which would differentiate the cloud market from other markets which have been subject to competition scrutiny: i) most customers (in revenue terms) are sophisticated enterprise customers, often employing IT specialists and using third-party IT consultants; ii) the cloud market is not a mature and stable market, i.e. there is still a large number of new customers starting to use cloud relative to the number of existing cloud customers;[252] iii) the cloud market is not a digital multi-sided platform market with (direct or indirect) network effects.

A3.110 In addition, Microsoft referred, to the extent that analogies with other markets are at all useful, to the recent market study into music streaming undertaken by the CMA, which was not referred for a market investigation reference. In particular, Microsoft referred to the finding by the CMA of positive outcomes for consumers (notably, a vast selection of music available at declining prices in real terms in recent years), notwithstanding limited switching and the existence of some barriers to switching (e.g. playlists curated by consumers over time or personalised recommendations made by the streaming services).

A3.111 Lastly, Microsoft was of the view that the perceived low level of switching in the market for cloud services is not caused by contractual restrictions imposed by cloud providers and referred to some findings of the Context Consulting market research commissioned by Ofcom, notably that switching is "primarily a function of internal factors rather than

[249] [✂].
[250] [✂].
[251] Microsoft response to the interim report.
[252] According to Microsoft, examples of mature markets (with a very low proportion of new customers entering the market for the first time compared to the number of already existing customers) would be: retail banking, mobile telephony and broadband Internet.

provider-imposed restrictions."[253] In fact, according to Microsoft, the prevalence of switching in the UK – as measured by Context Consulting[254] – would be consistent with a high level of competition.

A3.112 After considering Microsoft's views on switching, our view is that they do not alter our conclusion that there is a low level of switching, as we explain below and elsewhere in this report.

A3.113 We acknowledge that, as a general proposition, low levels of switching in a market may not always be evidence of weak competition. However, as we have explained in Section 5, we have found that there are several barriers to switching between providers of cloud infrastructure services. We also consider that general findings from other markets – such as those mentioned by Microsoft (music streaming, retail banking, mobile telephony and broadband internet) – have limited bearing on our evidence and market-specific analysis of the cloud services market. Each market study looks at the specific facts and characteristics of the market under consideration. Regarding network effects, we have found that they do exist in the cloud market, as discussed in Section 6.

A3.114 Lastly, we consider that Microsoft's references to the findings of the Context Consulting market research do not properly reflect the scope and findings of that research. In particular, Context Consulting noted that it had "encountered few if any examples of organisations switching away from one of the hyperscalers" and that it "is still relatively early in the adoption journey for most companies, and they are evaluating progress rather than looking to make significant changes. In most cases, firms are still on the way in, not out, of their IaaS / PaaS environments. Switching tends to be done with a relatively small portion of data and workloads, moving from one minority provider to another."[255] Furthermore, Context Consulting found that "most decision-makers acknowledge that a de facto lock-in exists," although they also noted (as Microsoft pointed out) that "this is primarily a function of internal factors rather than provider-imposed restrictions."[256]

Evidence from [✂] on absence of lock-in – and our assessment

A3.115 In response to our interim report,[257] [✂] submitted a quantitative analysis based on all opportunities in which [✂] had participated globally between [✂] and had a value higher than $[✂] million. [✂] made two substantive points. First, according to [✂], the increasing

[253] Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, pages 110 and 121.
[254] Context Consulting found that 18% of respondents have actively switched providers for all existing workloads; 35% of respondents have switched to another provider for new workloads; another 35% of respondents are considering switching provider; 23% of respondents have never considered switching. See Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, page 105.
[255] Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, page 110.
[256] Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, page 110.
[257] [✂]. In its earlier submission, [✂]. According to [✂], this was evidence of absence of lock-in. [✂] also noted [✂]. As we note in the text, we think that the methodology used by [✂] to classify customers as new or existing in the subsequent submission still presents significant shortcomings.

number of opportunities issued by its existing customers globally shows that customers did not consider themselves to be "locked-in" with [✂].

A3.116 Second, [✂] calculated its win rate in tenders[258] issued by existing customers globally and compared it with the win rate in tenders issued by new customers globally. According to [✂], this analysis shows that its win rate in tenders from existing customers is [✂] its win rate in tenders from new customers and that both rates are [✂] over time. According to [✂], the "[✂] existing customers are naturally more willing to continue to use [✂] services given their positive customer experience. Importantly, the results show [✂], which demonstrates that [✂] win rate in tenders is not driven by existing customers being locked in. Overall, [✂], which suggests increased competitive pressure from other cloud providers."[259]

A3.117 As we explain below, however, we have reservations regarding the approach and the data used by [✂] for its quantitative analysis. We also disagree with [✂] about its two substantive findings.

A3.118 We note that [✂] excluded from its analysis all opportunities with an estimated value below $[✂] million. By way of illustration of the potential impact of this filtering out by [✂], [✂] told us that, for [✂] (the most recent year of complete data), it included [✂] opportunities relating to UK customers in its analysis which were worth approximately $[✂] million in total. For the same year, [✂] left out [✂] opportunities relating to UK customers (as each of those opportunities was individually worth less than $[✂] million), worth approximately $[✂] million in total.[260] In other words, for every opportunity relating to UK customers included in its analysis, [✂] excluded roughly [✂] other opportunities relating to UK customers. Since [✂] win rate is calculated simply as the ratio between the number of opportunities won and the total number of opportunities it had participated in, leaving out such a high number of tenders is likely to have a material impact on [✂] findings.[261]

A3.119 Furthermore, despite [✂],[262] we think that [✂] may still not fully distinguish existing customers from new ones – as also acknowledged by [✂].[263] In particular, by way of example, by filtering out individual opportunities worth less than $[✂] million, [✂] may have classified customers who at a certain point in time had allocated an opportunity below the threshold and had then issued a subsequent tender for an opportunity above the threshold as new (rather than existing) customers.[264] We think that this inability by [✂] to fully and accurately classify customers as existing or new customers in its dataset has the potential to bias the results of its analysis.

A3.120 Lastly, and as also acknowledged by [✂][265] and noted in paragraph A3.48 above, we have reservations about the general quality, reliability, and comprehensiveness of the information

---

[258] As noted in footnote 197, the term "tender" used by [✂] is to be interpreted loosely. The dataset used by [✂] lists commercial opportunities, which may not always be formal tenders or negotiation processes.
[259] [✂].
[260] [✂].
[261] For clarity, the illustration above refers to opportunities relating to UK customers, whereas [✂] full results refer to global opportunities (despite [✂] apparent ability to extract data for tenders relating to UK customers only). [✂] has not provided an analysis of win rates including UK tenders only.
[262] See footnote 257.
[263] See [✂]. According to [✂], it "does not typically distinguish between opportunities that arose from new vs existing customers, and does not systematically track this type of information at the opportunity level in [✂]."
[264] As noted above, in its submission dated [✂], [✂].
[265] [✂] response dated [✂] to the s.174 notice dated [✂].

recorded in the dataset used for this analysis. Moreover, we note that the dataset used by [✂] does not capture, as an opportunity, [✂]. In addition, for this specific analysis, [✂] has used data for global (as opposed to UK) opportunities, which may not properly reflect [✂] win rate for opportunities from its UK customers. Accordingly, we question the suitability of this dataset to analyse the extent to which existing customers may face barriers to multi-cloud and switching, and how this may impact their ability to threaten to move workloads to [✂] competitors.

A3.121 On the first substantive finding by [✂] (i.e. that the number of opportunities issued globally by its existing customers has increased over time, which, according to [✂], would suggest absence of lock-in), we do not think that the increase in the number of opportunities in the data is relevant to analyse the extent to which customers may face barriers to multi-cloud and switching. We also note that the increase in the number of opportunities is largely driven by an increase in the number of customers present in [✂] dataset over time. In addition, the growth in additional opportunities is likely to reflect the fact that existing customers take up additional services over time (either because they move more workloads to the cloud or new use cases arise). While customers may have been willing to explore services offered by [✂] competitors, the fact that there are more opportunities in the dataset does not imply that customers ultimately considered these rivals as viable alternatives.

A3.122 On the second substantive point raised by [✂] (i.e. that its win rate for existing customers is [✂] for new customers), we note that – even on the evidence submitted by [✂] and putting to one side the concerns with that data we raise above – the data shows a [✂] and that [✂] wins [✂] for existing customers. In our view, this is consistent with some existing customers finding it difficult to switch existing workloads and allocate new workloads to cloud providers other than [✂], even if [✂] win rate [✂] and the difference between win rates for existing and new customers [✂] over time.

A3.123 Overall, we conclude that [✂] views do not alter our assessment about the extent of lock-in the cloud infrastructure services market.

A4. Further detail on technical barriers to multi-cloud and switching

A4.1 In this annex we set out further detail on the technical barriers to multi-cloud and switching that we have identified in Section 5. To do so, the annex is structured as follows:

a) First, we provide further details about the concepts of interoperability and portability in the cloud.
b) Second, we set out further evidence and details regarding technical barriers to different types of multi-cloud
c) Third, we set out further evidence and details regarding technical barriers to different types of switching.
d) Fourth, we set out data provided by hyperscalers indicating that a significant portion of customers is likely to face material barriers to multi-cloud and switching.
e) Finally, we present a more detailed overview of AWS arguments in relation to asymmetry of functionalities.

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

A4.2 As discussed in Section 5, customers rely on cloud infrastructure services to run applications and process data which they use to provide services internally and/or externally to their users. To ensure their cloud architectures work as effectively as possible, customers typically face three categories of technical challenges: interoperability, application portability and data portability challenges. We discuss these challenges 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 A4.1: Stylised example of customer application relying on a cloud service

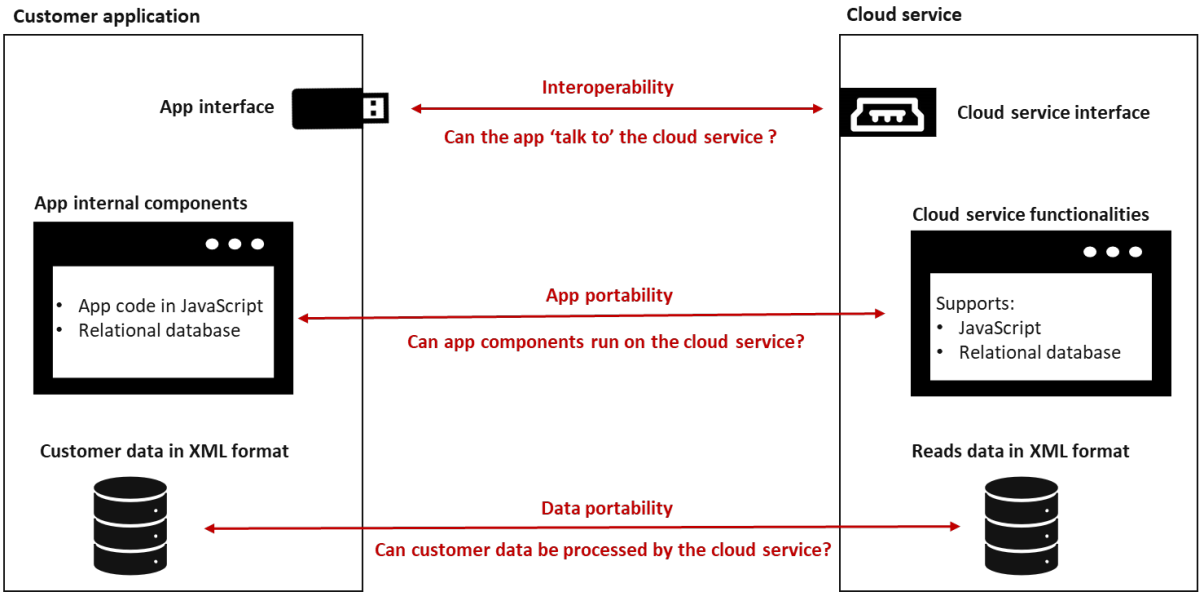

Source: Ofcom.

A4.3 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.[266] Similarly, interoperability is a key consideration if a customer wants to integrate different cloud services hosted on several clouds.[267] To the extent the two cloud providers use incompatible interoperability rules, the customer may need to develop or procure an adaptor[268] 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 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.[269]

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 in cloud infrastructure services compared to interoperability and application portability.[270]

---

[266] 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 Figure A4.1, 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.

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

[268] 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.

[269] 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.

[270] 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 in cloud infrastructure services.

A4.4 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.

A4.5 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.

A4.6 As set out in Section 5, we consider that there are several factors which can reduce interoperability and portability in cloud and, as a result, hinder customers' ability to implement different multi-cloud or multi-vendor architectures and switch. These factors include: technical differentiation of cloud infrastructure services, technical differentiation of ancillary cloud services, integration and operationalisation efforts, asymmetry of functionalities, data gravity, lack of technical skills and lack of transparency.

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

A4.7 This subsection considers the extent to which a lack of interoperability and portability may hinder customers' ability to implement the following multi-cloud architectures:[271]

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.

A4.8 As noted in Section 5, the evidence we have received since the interim report has clarified that there is not a clear-cut separation between different multi-cloud architectures. Instead, there is a spectrum of possible implementations going from the ones with little to no integration (i.e. the more 'silo-ed' end of the spectrum) to the more highly integrated ones (i.e. the more 'integrated' end of the spectrum).[272]

[271] 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.
[272] An extreme version of integrated multi-cloud is where applications dynamically distribute microservices between different clouds depending on a predetermined set of rules (e.g. which cloud offers the cheapest storage).

A4.9 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 A4.2: Overview of findings on the technical effort to implement multi-cloud

Multi-cloud scenario	Technical barriers	Customers most affected by barriers	Level of effort
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
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
Siloed multi-cloud	Costs of maintaining parallel clouds.	Customers with a cloud architecture tightly integrated with proprietary services of cloud providers.	Medium

Source: Ofcom.

Interoperability and portability considerations when using multiple integrated clouds

A4.10 With integrated multi-cloud different customer applications, customer data and/or cloud services hosted on two or more public clouds are highly integrated.

A4.11 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.

A4.12 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).[273] With respect to cloud services, even where they can 'talk to each other', services from

[273] 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).

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.[274]

A4.13 As discussed in Section 5, 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.

A4.14 Overall, as discussed in Section 5, we consider that technical barriers to deploy an integrated multi-cloud architecture are likely to be material and an increasing function of the desired level of integration between customers applications, data and cloud services hosted on different clouds.

Box A4.3: Hypothetical example of customer deploying integrated multi-cloud

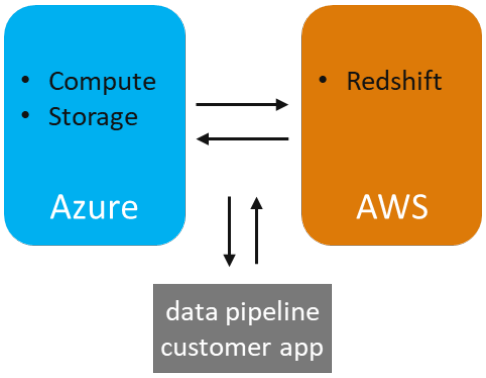

A digital advertising agency is running its main operation in Microsoft Azure where they use compute, storage and various other cloud services to deliver customised advertisements to their customers. When processing the resulting data to assess performance, they want to use the Amazon RedShift data warehouse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

Interoperability and portability considerations when duplicating clouds

A4.15 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.

A4.16 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

[274] Customers will face additional 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.

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).[275]

A4.17 As noted in Annex 3, 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 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.[276] 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.[277, 278]
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, [✂] 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.[279]

A4.18 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.[280]
b) Large organisations responding to our customer questionnaire indicated cost was a key barrier. For example, [✂] explained that it would be cost prohibitive to maintain both

[275] 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 Gov.uk. Managing technical lock-in in the cloud [accessed 29 September 2023]). This is also confirmed by evidence we have received from stakeholders. For example, [✂] 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 / [✂] meeting, [✂]); [✂] said that external facing workloads would have a significant technical barrier as they utilize more cloud native services ([✂] response dated [✂] to the s.174 notice dated [✂], question [✂]); Ofcom / [✂] meeting, [✂]; Ofcom / [✂] meeting, [✂]; [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[276] Context Consulting research data tables, Q29. This figure is 14% after excluding respondents using only private cloud.
[277] More specifically, 38% of respondents mentioned lack of interoperability between clouds, and 36% of respondents mentioned technological difficulties. Context Consulting research data tables, Q31.
[278] 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 3, it is difficult to assign these responses to a specific multi-cloud architecture.
[279] [✂].
[280] Analysis of Context Consulting research data tables, Q31.

clouds actively all the time. Therefore, any multi-cloud alternative would need to be limited but with the ability to scale up rapidly.[281] [✂] indicated replication of services as the most important obstacle for multi-could as it drives cost and complexity of the overall environment.[282] 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.[283]

A4.19 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. 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.

A4.20 Overall, we 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 A4.4: Hypothetical example of customer duplicating clouds for resilience

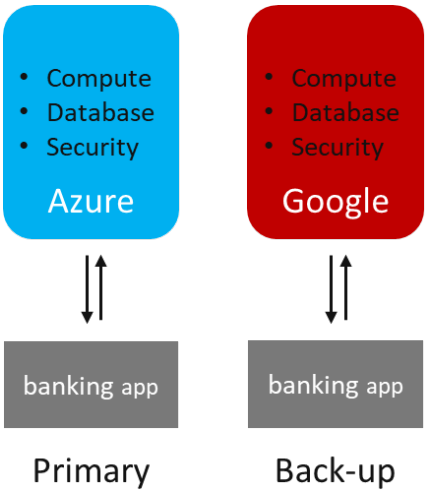

A small online payment provider wants to ensure that customers are still able to make payments during an outage. To minimise cost and complexity, they decide to replicate a part of their existing architecture to a different cloud provider to improve reliability and resilience for the most essential services. 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.

[281] [✂].
[282] [✂] response dated [✂] to our customer questionnaire, question [✂].
[283] BT Group response to the CFI, page 9.

Interoperability and portability considerations when using multiple siloed clouds

A4.21 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).

A4.22 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.[284] This technical effort would be more acute where such services have a low level of standardisation.

A4.23 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.

A4.24 As discussed in Section 5, the evidence received indicates that, where it occurs, multi-cloud is predominantly siloed as this is less technically demanding.

A4.25 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, [✂] 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, such as previous deployments and skills sets, rather than distinct workloads for specific cloud providers.[285]

A4.26 Overall, our 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 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.

A4.27 As discussed in Section 5, 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.

---

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

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

Box A4.5: Hypothetical example of a customer deploying siloed multi-cloud

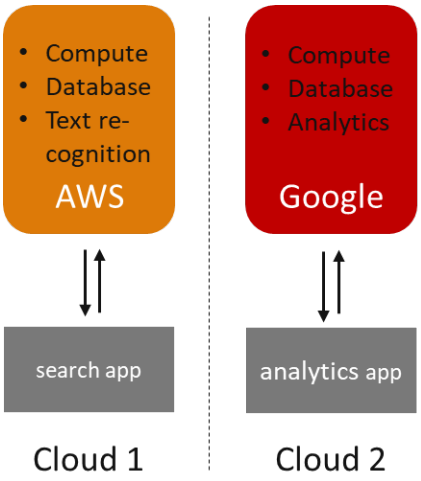

A large telecommunications vendor has multiple organisational divisions with only little overlap in their activities which requires no integration between their services. Each division has their own dedicated technical resources and significant autonomy when it comes to architectural infrastructure decisions 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

A4.28 In this subsection we consider the extent to which a lack of interoperability and portability may hinder customers' ability to switch:

a) **Between public clouds**: customers switch from using cloud A (origin cloud) to using cloud B (target cloud). This can be done for the entirety of the customer architecture (full switch) or for parts of it where only certain existing workloads are migrated (partial switch).
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).

A4.29 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 A4.6: Overview of findings on the technical effort to switch

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

Source: Ofcom.

Interoperability and portability considerations when switching between public clouds

A4.30 Switching between clouds involves customers recreating their cloud architecture on the target cloud so that they can cease using the origin cloud.

A4.31 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.

A4.32 As discussed in Section 5, the evidence we have received, and our own analysis, indicates that technical challenges are likely to be material for some customers and use-cases. In addition to the evidence set out in Section 5, we note that:

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 to the market research, application portability was the third most frequently cited barrier to switching (31%), followed by data portability in fourth

place (30%) and interoperability challenges in eighth place (21%).[286] 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%)[287] and may also reflect the time and cost needed to reconfigure data and applications.

A4.33 As discussed in Section 5, the evidence received also indicates that more mature companies and larger companies may be more affected by technical barriers to switching. This may be because these companies are more likely to have large numbers of applications and/or use various proprietary services offered by their cloud providers. In addition to the evidence set out in Section 5, we note that:

a) 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.[288] This may be because more mature customers have more cloud applications.
b) Our 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.[289] 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.[290] Also, see discussion of average number of cloud infrastructure services that customers use with each hyperscaler in Section 5.

A4.34 Overall, we 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

---

[286] 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%).
[287] Context Consulting research data tables, Q52. This figure is 43% after excluding respondents using only private cloud.
[288] Analysis of Context Consulting research data tables, 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).
[289] Analysis of Context Consulting research data tables, 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 different 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.
[290] Context Consulting research report, slide 37.

applications or data 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.

Box A4.7: Hypothetical examples of customers switching public clouds

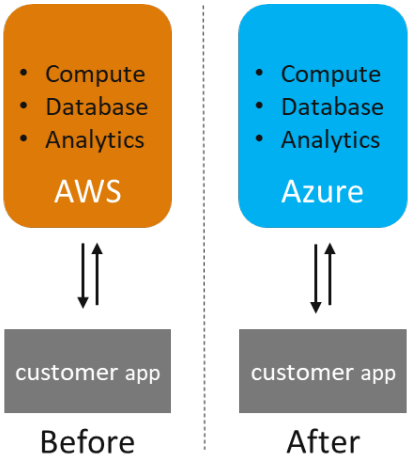

Switching the cloud provider of a single **blog application** which relies on a commonly used software package or stack may be an easy endeavour and may be completed within a day or two by qualified personnel.

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

A4.35 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.

A4.36 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.

A4.37 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.

A4.38 In line with the above, the evidence we have received, indicates that barriers to this type of switching may be lower compared to a switch between clouds:

a) One ISV ([✂]) 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.[291] Nonetheless, as outlined below, some ISVs ([✂]) said that switching between a proprietary service to an equivalent service offered by a third party would still require customers to rewrite a significant part of their code.
b) We have also received some feedback from customers that is consistent with the above. Specifically, an AWS customer ([✂]) 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.

A4.39 Overall, we 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.[292]

[291] [✂] response dated [✂] to the s.174 notice dated [✂], Part A question [✂].
[292] Ofcom / [✂] meeting, [✂]; and Ofcom / [✂] meeting, [✂].

Box A4.8: Hypothetical example of a customer switching within a public cloud

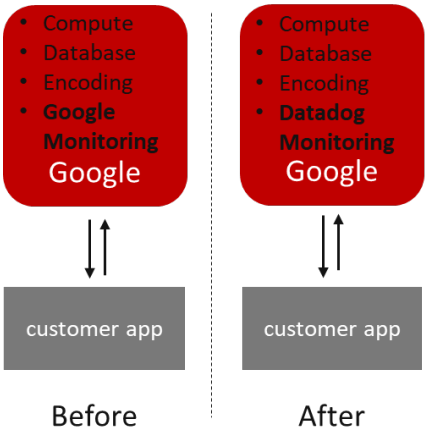

In this scenario an online communications network is looking to switch their monitoring system from a first-party solution to a third-party product, provided by an ISV.

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.

Data provided by hyperscalers indicates that a significant portion of customers is likely to face material barriers to multi-cloud and switching

A4.40 To assess the share of customers that may be facing high technical barriers to switching we asked hyperscalers to provide us with the number of first-party proprietary services and of first-party PaaS services[293] used by customers of different sizes. While we recognise there is not a perfect correlation between the number of services used and the degree of technical barriers faced, this analysis indicates that the number of proprietary services (and particularly the number of PaaS services) used by customers is likely to be indicative of the level of technical efforts customers would face if they wanted to switch or replicate some or all of their architecture across multiple clouds.

A4.41 As shown in the tables below, the responses indicate that, across all hyperscalers, customers spending more than £10k per year (which represent a significant portion of the overall hyperscaler's revenues[294]) consume at least [✂] [10-20] first-party proprietary cloud services and at least [✂] [5-15] PaaS services. The number of first-party proprietary services is generally higher for customers in higher revenue bands. For example, customers spending more than $1m a year (accounting for [✂]% of hyperscalers' revenue) on average take at least [✂] [30-40] first-party proprietary services and more than [✂] [20-30] PaaS services.

A4.42 This analysis indicates that the significant portion of customers taking [✂] [10-20] or more proprietary services would likely face a high degree of technical complexity if they wanted to switch or replicate some or all of their architectures across multiple clouds. In addition, the

[293] More precisely, we asked hyperscalers to provide the 'Average number of first-party public cloud infrastructure services used' as well as the 'Average number of first-party public compute, storage and networking services used'. To calculate the average number of PaaS services we subtracted the latter from the former.

[294] Specifically, for all hyperscalers, these customers account for [✂]% of revenues.

data also indicates that more mature cloud customers are likely to face more material technical efforts, which is consistent with the results from the market research.

Table A4.9: Average number of cloud infrastructure services by customer revenue band and cloud provider

	Average number of first-party public cloud infrastructure services used		
Annual revenue band ($m)	**AWS**	**Microsoft**	**Google**
<10k	[✂]	[✂]	[✂]
10k-1M	[✂]	[✂]	[✂]
1M-5M	[✂]	[✂]	[✂]
5M-10M	[✂]	[✂]	[✂]
10M-20M	[✂]	[✂]	[✂]
> 20M	[✂]	[✂]	[✂]

Source: Ofcom analysis of data provided by AWS,[295] Microsoft[296] and Google.[297]

Table A4.10: Average number of PaaS services by customer revenue band and cloud provider

	Average number of first-party public cloud infrastructure services used – excluding compute storage and networking services		
Annual revenue band ($m)	**AWS**	**Microsoft**	**Google**
<10k	[✂]	[✂]	[✂]
10k-1M	[✂]	[✂]	[✂]
1M-5M	[✂]	[✂]	[✂]
5M-10M	[✂]	[✂]	[✂]
10M-20M	[✂]	[✂]	[✂]
> 20M	[✂]	[✂]	[✂]

Source: Ofcom analysis of data provided by AWS,[298] Microsoft[299] and Google.[300]

[295] AWS response dated 14 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' column E).
[296] Microsoft response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Confidential Annex 2, tab 'Q1 Customer distribution' column E).
[297] Google response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' column E).
[298] These figures were calculated by subtracting the 'Average number of first-party public compute, storage and networking services used' from the 'Average number of first-party public cloud infrastructure services used'. AWS response dated 14 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' columns E, F).
[299] These figures were calculated by subtracting the 'Average number of first-party public compute, storage and networking services used' from the 'Average number of first-party public cloud infrastructure services used'. Microsoft response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Confidential Annex 2, tab 'Q1 Customer distribution' columns E, F).
[300] These figures were calculated by subtracting the 'Average number of first-party public compute, storage and networking services used' from the 'Average number of first-party public cloud infrastructure services used'. Google response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' columns E, F).

Table A4.11: revenue share of customers in different revenue bands

	Customer group spend as percentage of total customer spend		
Annual revenue band ($m)	**AWS**	**Microsoft**	**Google**
<10k	[✂]%	[✂]%	[✂]%
10k-1M	[✂]%	[✂]%	[✂]%
1M-5M	[✂]%	[✂]%	[✂]%
5M-10M	[✂]%	[✂]%	[✂]%
10M-20M	[✂]%	[✂]%	[✂]%
> 20M	[✂]%	[✂]%	[✂]%

Source: Ofcom analysis of data provided by AWS,[301] Microsoft[302] and Google.[303]

AWS and Microsoft disagreed with our view that technical barriers are higher if first-party cloud services are not fully functional when used in combination with third-party cloud services

A4.43 As discussed in Section 5, AWS and Microsoft disagreed with our preliminary view that technical barriers are higher if first-party cloud services are not fully functional when used in combination with third-party cloud services.

A4.44 This section provides further details regarding AWS submissions on this. Microsoft submissions are instead fully discussed in Section 5.

AWS submission

A4.45 As discussed in Section 5, AWS noted that the interim report only identifies interoperability limitations in a small subset of its cloud services, focusing on 10 services for which there are competing software solutions available that customers can run on AWS (or elsewhere). In addition, AWS said that the features identified as limiting interoperability are described inaccurately, exist alongside features that ensure interoperability, or are the product of an objective technical limitation.[304]

[301] To calculate these percentages the 'Revenue for public cloud infrastructure services' associated to each Annual revenue band was divided by the sum of 'Revenue for public cloud infrastructure services' across all Annual revenue bands. AWS response dated 14 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' column D).
[302] To calculate these percentages the 'Revenue for public cloud infrastructure services' associated to each Annual revenue band was divided by the sum of 'Revenue for public cloud infrastructure services' across all Annual revenue bands. Microsoft response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Confidential Annex 2, tab 'Q1 Customer distribution' column D).
[303] To calculate these percentages the 'Revenue for public cloud infrastructure services' associated to each Annual revenue band was divided by the sum of 'Revenue for public cloud infrastructure services' across all Annual revenue bands. Google response dated 11 July 2023 to the s.174 notice dated 13 June 2023, question 1 (Annex 2, tab 'Q1 Customer distribution' column D).
[304] AWS response to the interim report, paragraphs 24-26.

A4.46 AWS provided specific comments on each of the 10 cloud services identified in the interim report.[305] However, we consider that AWS has not provided sufficient evidence to rebut our analysis and our view that technical barriers are higher if first-party cloud services are not fully functional when used in combination with third-party cloud service.

A4.47 In the following paragraphs, for each of the 10 services listed in Table 5.9 of the interim report, we set out: (i) the technical limitation we had identified at the interim report stage; (ii) the AWS rebuttal to our analysis; and (iii) our views based on AWS submission.

Amazon SageMaker and SageMaker DataWrangler

A4.48 In our interim report we indicated that Amazon SageMaker can only access training data from Amazon S3, Amazon EFS and Amazon FSx.

A4.49 AWS disagreed that Amazon SageMaker can only access training data from certain AWS services. AWS acknowledged that Amazon SageMaker can only manage integrated resources (in this case, Amazon S3, Amazon EFS or Amazon FSx), but said that this does not limit interoperability as customers can instruct Amazon SageMaker to pull data from any data source, including other AWS services or third-party solutions (including other cloud providers or on-premises storage). AWS noted that it has published a blog post with instructions for downloading data directly from third-party data sources into a SageMaker Training job instance, and that this provides customers with the flexibility to bring in data from a source of their choice and in a format that they see fit for their use cases. AWS said that SageMaker DataWrangler, which is an optional feature of Amazon SageMaker, can query and import data from dozens of data sources, including other AWS services, Databricks, Snowflake, and more than 40 other third-party storage solutions.

A4.50 AWS said it enables customers to build the machine learning solution that works best for them by choosing the features they want to use as part of their machine learning lifecycle. Customers are free to use a single SageMaker feature for a discrete use-case (for example using SageMaker Studio to train models that will be deployed elsewhere), or use multiple SageMaker features for a more holistic solution. According to AWS, the requirements and dependencies of each feature, including their interactions with other AWS services, are clearly laid out in the Developer Guide. AWS said that Amazon SageMaker further enables portability and interoperability by using standard Docker containers for core ML tasks such as code authoring, training and hosting. Customers can bring their own Docker containers and use those containers for model building, training, and hosting. With minimal changes, customers are able to reuse the same Docker containers outside of Amazon SageMaker, both on-premise or with a different cloud provider. Amazon SageMaker supports popular open-source resources such as PyTorch and TensorFlow, and models developed on Amazon SageMaker using these frameworks can be deployed elsewhere without any changes.

A4.51 In addition, AWS said that customers can use an Amazon SageMaker SDK to integrate with other systems and services. Customer who have their ML systems built using Kubernetes clusters, self-managed or managed by another provider, can also integrate with Amazon SageMaker by using Kubernetes operators that SageMaker offers. This provides flexibility in switching services or having a multi-cloud architecture.

[305] [✂].

A4.52 We note that AWS's response points at a recently introduced feature (described in a blog post)[306] which allows customers to download data directly from third-party data sources into Amazon SageMaker. The new addition remedies the concerns we raised in our interim report and removes the limitation we mentioned in relation to Amazon SageMaker. As discussed in Section 5, we welcome the lifting of such restriction. However, the speed with which features are made available on third-party services may be a relevant factor for customers – in general, customers may be discouraged from creating a multi-cloud or multi-vendor architecture if they anticipate that the third-party cloud services they wish to use would gain access to the full set of functionalities of any relevant first-party cloud services only after some time.

Amazon Redshift

A4.53 In our interim report we indicated that Amazon RedShift can only bulk load data from Amazon S3.

A4.54 AWS claimed that there are many methods for loading data into Redshift, of which bulk loading is only one. AWS said that all of these methods, except for bulk-loading, support ingestion from any source, such as other cloud providers or on-premises data/file storage, via AWS Data Pipeline or SSH, through industry-standard JDBC/ODBC driver interfaces. Amazon Redshift also provides a Data API that allows users to access data from Amazon Redshift with any application, including on-premises, cloud-native, and containerised services.

A4.55 AWS said that bulk-loading data to Amazon Redshift is currently limited to Amazon S3 as the source because it is a technically complex task whereby data is uploaded in parallel to each compute node to maximize the rate at which data is ingested into the data warehouse cluster. According to AWS, bulk-loading is not a feature that can be easily rolled out across multiple services, as significant investment and customisation is required to enable any two services to coordinate data upload across parallel nodes. AWS noted that this technical challenge of getting the two services to coordinate data upload also applies to transfers between other AWS services and Redshift, which is why this bulk load option is not available between them. Instead, those other AWS storage and database services upload data to Amazon Redshift in the same manner as users of third-party services.

A4.56 We note that AWS acknowledged that Amazon RedShift can only bulk-load data from Amazon S3. Hence, we consider the technical limitation identified is still valid.

A4.57 We disagree with AWS that the limitations identified in relation to Amazon Redshift are not important because they exist alongside other features that ensure interoperability. In particular, while other data ingestion options may exist, none of them are comparable to the speed and efficiency of bulk-loading data. Bulk-upload is also a commonly used feature: large data sets are routinely loaded into data warehouses. Those processes can be essential for time critical applications.

A4.58 We also disagree with AWS that the limitations identified in relation to Amazon Redshift are just the product of an objective technical limitation. Our analysis suggests that, while some work may be required, it is not technically difficult to provide access to the bulk-upload

[306] AWS website. Use Snowflake as a data source to train ML models with Amazon SageMaker [accessed 12 September 2023]

functionality to third-party services. This outcome may be even easier to accomplish if all functionalities were made accessible from the design stage (i.e. 'interoperability by design').

Amazon Athena

A4.59 In our interim report we indicated that Amazon Athena can only query data stored on Amazon S3.

A4.60 AWS said that Amazon Athena is a serverless, interactive analytics service built on open-source, supporting open-table file formats, and allows customers to analyse petabytes of data. AWS stated that customers can analyse data or build applications from Amazon S3 or from dozens of other sources, including on-premises or other cloud providers, using available data source connectors. AWS said customers can also use the open-source Athena Query Federation SDK to build custom connectors to query any data source.

A4.61 We accept that customers can use available connectors or build their own connector to allow Athena to query data stored on other clouds. However, this would still involve additional technical work and overheads compared to the native option available for Amazon Athena. Moreover, even with a connector in place a customer would still need to copy all the data to S3 which will result in extra engineering, operational steps, storage cost and data duplication.

Amazon Interactive Video Service (Amazon IVS)

A4.62 In our interim report we indicated that Amazon IVS can only auto-record to Amazon S3.

A4.63 AWS claimed that auto-record to S3 is an optional feature of Amazon IVS that allows live videos to be saved to Amazon S3 for archiving and playback. Amazon IVS uses S3 for this feature because customers value S3's security and availability profile as an AWS service. AWS said that "auto-record" does not limit interoperability as customers have many ways to store their video streams. As Amazon IVS uses industry standard video formats, a customer who chooses to use the auto-record feature can easily move the recorded video to other cloud providers or on-premises storage or playback services. Alternatively, customers can choose to record video streams without using the auto-record to S3 feature at all. For example, customers can record video via their own or third-party client-side broadcasting services or server-side recording services, and store video in any location. Amazon IVS supports the industry standard RTMP streaming protocol, meaning customers can use Amazon IVS in conjunction with other services that they build or use.

A4.64 We note that AWS acknowledged that Amazon IVS can only auto-record to Amazon S3. Hence, we consider the technical limitation identified at interim report stage is still valid.

A4.65 We disagree with AWS that this limitation is not important since it exists alongside other features that ensure interoperability. We consider that limitations to effective interoperability may still exist, if services are not fully functional when used in combination with third-party services. For example, not enabling auto-record may add costs to using Amazon IVS with third-party services and discourage the take-up of multi-vendor or multi-cloud architecture.

Amazon Kinesis Video Streams

A4.66 In our interim report we indicated that Amazon Kinesis Video Streams can only deliver extracted images to Amazon S3.

A4.67 AWS claimed that, while it is the case that the Amazon S3 delivery feature extracts images from videos stored in Kinesis Video Streams and delivers those images to S3, this does not

limit interoperability. AWS said that customers can set up near real-time image extraction and delivery to any storage service –on-premises or on any other cloud provider – by using Kinesis Video Streams Notifications along with the GetImage API. Kinesis Video Streams stores data in an open-standard file type (MKV), which enables users to easily extract and process their data using third-party services, or move it to other storage services.

A4.68 AWS said that, as Kinesis Video Streams is a fully-managed service, customers use it specifically to ingest video data and store and process it on the AWS cloud. While the S3 delivery feature was designed based on known customer use cases, AWS has not received any customer requests to extract images from Kinesis Video Streams to a non-S3 destination, and has therefore not devoted resources to developing such a feature.

A4.69 We note that AWS acknowledged that Amazon Kinesis Video Streams can only deliver extracted images to Amazon S3. Hence, we consider the technical limitation identified at interim report stage is still valid.

A4.70 As discussed above, we consider that limitations to effective interoperability may still exist, if services are not fully functional when used in combination with third-party services. For example, with automated image extraction, images will be extracted from video based on customer-defined rules and no other intervention is necessary. With "manual mode", a customer will need to engineer and operate their own application to extract image data in regular intervals which is more costly, prone to errors, creates overheads and comes with a time delay. This would discourage the take-up of multi-vendor or multi-cloud architectures.

Amazon Omics

A4.71 In our interim report we indicated that Amazon Omics can only use Amazon S3 for data import and export

A4.72 AWS said that Amazon Omics is a new AWS service which was publicly launched in December 2022 and new features continue to be developed, including features that further assist customers in integrating Amazon Omics with other IT solutions. AWS said it recently released a feature enabling Omics Storage to receive inputs from, and export data to, on-premises or local storage in addition to S3. AWS said features enabling further interoperability are currently being developed and will be released in the future.

A4.73 We acknowledge that AWS has introduced new features which allow customers to receive inputs from, and export data to, on-premises or local storage in addition to S3. However, we note that Amazon Omics still does not support other clouds storage services for data import and export. Hence, we consider the technical limitation identified at interim report stage is still valid.

Amazon Pinpoint

A4.74 In our interim report we noted that Amazon Pinpoint does not allow adding Amazon Personalise recommendations to a marketing email campaigns if these are from third-party recommendations engines.

A4.75 AWS stated that it is not the case that customers wishing to use Amazon Pinpoint can add Amazon Personalize recommendations to a marketing email campaign but cannot add recommendations from third-party engines. AWS said that customers are also able to leverage personalised recommendations generated by third-party services by updating the attributes of endpoints stored in Pinpoint using the UpdateEndpoint API. According to AWS, many customers are using third-party recommendations tools in this way today.

A4.76 We understand that the Amazon Pinpoint API currently only supports Recommender Models that require customers to deploy an Amazon Personalize campaign[307] (i.e. the limitation set out in our interim report). We accept that customers can update recommendations manually to incorporate inputs from third-party services (i.e. what AWS refer to as updating the attributes of endpoints stored in Pinpoint using the UpdateEndpoint API). However, this can be regarded as work-around to existing limitations, rather than a solution that promotes interoperability.

Amazon Timestream

A4.77 In our interim report we noted that Amazon Timestream can only use AWS Backup for backups.

A4.78 AWS said that Amazon Timestream interoperates with customers' preferred data collection, visualization, analytics, and machine learning services where they are hosted, allowing customers to include Amazon Timestream in their time series solutions. Customers can connect to Amazon Timestream through popular third-party tools such as open source Telegraf, Apache Flink, and a variety of AWS services, and visualize their data through tools such as Grafana or any other source using standard JDBC connectors.

A4.79 AWS acknowledged that customers wishing to use Amazon Timestream can only use AWS's Backup service to manage backups. It claimed this is not a limit on interoperability as customers can use the JDBC connector to extract their data out of Timestream and back it up in a location of their choice using the backup service of their choice, whether third-party, another cloud provider or on-premises.

A4.80 We note that AWS acknowledged that Amazon Timestream can only use AWS Backup for backups. Hence, we consider the technical limitation identified at interim report stage is still valid.

A4.81 As discussed above, we consider that limitations to effective interoperability may still exist, if services are not fully functional when used in combination with third-party services. For example, we accept that customers can use available connectors to back-up outside of AWS. However, such an option would involve addition technical work and overheads compared to the native option available for Amazon Timestream.[308]

A4.82 In addition, we note that AWS recently enabled a new unload functionality for Amazon Timestream which, based on publicly available documentation,[309] would appear to only work with Amazon S3.

AWS IoT Events

A4.83 In our interim report we noted that AWS IoT Events can only trigger actions with other AWS services.

A4.84 AWS disagreed that customers wishing to use AWS IoT Events can only trigger actions with other AWS services. AWS said customers can use AWS IoT Events to trigger actions in any service – on another cloud provider or on premises – using an AWS Lambda function. Lambda can be configured to receive an event document from IoT Events, convert it into a set of instructions (an object), and then pass the object to any pre-defined destination using

---

[307] AWS website. Recommender Models [accessed 29 September 2023].
[308] We also note that JDBC connectivity is available only to customers utilising Java and other JVM-based languages.
[309] AWS website. Using UNLOAD to export query results to S3 from Timestream [accessed 29 September 2023].

APIs. Customers can define objects in several common languages such us Java, Go, PowerShell, Node.js, C#, Python, and Ruby, enabling easy communication with other applications.

A4.85 We understand that to allow AWS IoT to become interoperable with non-AWS services, a customer is currently required to purchase another AWS service (AWS Lambda) to engineer their own bridge solution. While we accept this may be possible, such an option is likely to require additional technical efforts on the customer side. Moreover, the fact that a bridge solution can only be built using another proprietary AWS cloud service (AWS Lambda) corroborates our view that AWS IoT Events can only be used in combination with other AWS products.

A4.86 Lastly, we note that the list of services identified in Table 5.2 of this document is not meant to be an exhaustive list of all AWS (or Microsoft) services featuring these limitations but rather a list of examples to better illustrate what 'asymmetry of functionality' could look like in practice.

A5. Discounting analysis

A5.1 In response to the interim report, two hyperscalers disputed our market research findings that price rises are relatively common amongst renegotiating customers. They submitted evidence to argue that prices have not increased for customers.

A5.2 The evidence presented by the hyperscalers mainly focused on discount outcomes for cross-service privately negotiated committed spend customers. In this annex, we consider their evidence on discount outcomes and set out our assessment of their evidence.[310] We note that this annex does not assess each individual piece of evidence submitted by the hyperscalers but is an assessment of their key findings in the round, which includes our own observations.

A5.3 In summary, we acknowledge that the evidence presented by the hyperscalers does not support our market research findings which suggests that price rises are relatively common amongst renegotiating customers. However, we think that our research is likely to be capturing other issues. We also do not think the evidence received from the hyperscalers is sufficiently comprehensive, or robust enough, to allow us to fully test whether committed spend discount customers that face greater barriers to switching or multi-cloud may be subject to worse outcomes. Based on the data, we also observe significant variation in the discounts received by some customers, which suggests that the bargaining power of customers is likely to be important in negotiations, and that there is potential scope for worse outcomes for customers who face greater barriers to switching and multi-cloud.

A5.4 In the remainder of this annex, we first consider the evidence on discount outcomes for committed spend discount customers that renewed their contracts. We then consider evidence on the distribution of discounts and importance of customers bargaining power in negotiating discounts. Lastly, we outline the evidence received from customers on their experience of negotiating with the hyperscalers, and the extent to which bargaining power influences outcomes more generally.

Outcomes for committed spend discount customers

A5.5 In this subsection, we first summarise the main arguments and findings presented by the hyperscalers on discount outcomes for renegotiating committed spend discount customers. We then present our own assessment of their submissions, including why we do not think the data and analysis presented by the hyperscalers is able to fully test whether customers that face barriers to switching and multi-cloud are subject to worse outcomes.

Hyperscalers' submissions

A5.6 In response to our interim report, two hyperscalers [✂] and [✂] submitted some analysis on the privately negotiated committed spend discounts agreed with their customers.[311] Based on this analysis, the two hyperscalers find that, on average, these customers do not

[310] We also received some evidence from hyperscalers to illustrate how list and net prices have evolved over time. We separately assess evidence on list and net prices in Section 4.
[311] [✂]. [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂].

receive worse discounts upon renegotiation, once [✂]. They conclude from this analysis that [✂].

A5.7 In particular, one hyperscaler [✂] stated that they do not decrease the discounts available to customers at renegotiations. Instead, customers with new and renegotiated deals tend to, on average, receive [✂] discounts once [✂].[312] Based on their analysis, they find that:

i) On average, UK customers with new and renegotiated deals tend to receive a [✂] level of discount, once [✂].[313]
ii) Discounts [✂] for [✂]% of UK customers upon renegotiations (for renegotiated deals with an increase in committed spend, discounts [✂] for [✂]% of customers), and they estimate that discount rates [✂] on average by [✂]% for these customers.[314]
iii) There are some instances where customers receive [✂] discounts when committing to similar spend levels upon renegotiations, however this is due to the specific circumstances of the deal, [✂].[315]

A5.8 Another hyperscaler [✂] stated that discounts for committed spend customers do not decrease at contract renewal and if anything, discounts on average appear to be higher for some renegotiating customers at the global level, when [✂].[316] Based on their analysis, they note that:

i) On average, UK and global customers tend to receive similar or higher discounts when negotiating subsequent contracts.[317]
ii) Only a small number of customers receive lower discounts upon renegotiation of their contract, and that these cases [✂].[318]

A5.9 In addition, the hyperscaler, [✂], also stated that our market research findings on price rises could be driven by customers interpreting increases in their spend due to higher consumption, as price rises. They also argued that our question in the research asked if respondents had experienced a price rise for "some or all services", and therefore our findings on price rises could apply to only one service, rather than all services taken up by the customer.[319]

Our assessment of hyperscalers' submissions

A5.10 Our concern is that customers that face significant barriers to switching and multi-cloud have a limited ability to credibly threaten to switch or move some of their workloads to new providers and face worse terms than they would have been able to achieve otherwise.

A5.11 In the interim report, we referred to our market research finding that price increases are common for customers renegotiating, which could be an indication of harm to customers. The survey found that around 44% of respondents who renegotiated a contract reported

[312] [✂].
[313] [✂].
[314] [✂].
[315] [✂].
[316] [✂].
[317] [✂].
[318] [✂].
[319] [✂].

experiencing a price rise for some or all of their services, with an average price rise of 20% and a median price rise of 10% being cited.[320]

A5.12 Two hyperscalers, [✂] and [✂], have sought to test empirically whether these research findings are consistent with their data. We have reviewed the data provided by the hyperscalers on discount outcomes for committed spend discount customers that renegotiate their contract. Their analysis shows that for [✂], on average, committed spend discount customers do not receive worse discounts upon renegotiations. However, we have some concerns around the relevance and robustness of their analysis, which we discuss further below. Based on this, we do not think their analysis is able to conclusively test whether customers that face greater barriers to switching or multi-cloud may be subject to worse outcomes.

A5.13 We also observe that a [✂] of customers do receive [✂] upon renegotiations, [✂].[321] However, the hyperscalers suggest that this is not necessarily indicative of worse outcomes for customers. One hyperscaler, [✂], explained that this may reflect customers that are [✂],[322] whilst [✂], explained that lower discounts are typically given due to "specific business reasons".[323]

A5.14 Nevertheless, we acknowledge that our market research findings on price rises does not support the analysis presented by the hyperscalers. We think the price rises captured as part of our research could be reflecting other changes or issues, and that the findings may be more reflective of pricing outcomes for customers without a committed spend discount.

i) Our market research suggests that only 6% of respondents have a committed minimum spend in place.[324] Therefore, our research findings on price rises are unlikely to be reflective of outcomes for committed spend discount customers, which is the core focus of [✂] and [✂] analysis.

ii) Our market research also asked respondents if they had experienced a price rise for "some or all services". Our findings could therefore also be reflective of those customers that may have experienced price rises for some individual services only. For example, in Section 4, we noted that [✂].[325] In recent times, one hyperscaler, [✂], which led to a price increase for [✂] products in the UK.[326] On the other hand, we also recognise that our findings could be driven by some customers mis-interpreting our question and reporting about increases in spend experienced as a result of increasing their own usage with a given hyperscaler.

iii) We also think these price rises could be reflective of other issues. For example, bill shock arising from some customers finding it difficult to predict their cloud spend; the take up of ancillary services (e.g. [✂]) which can represent an effective

---

[320] Context Consulting research report, slides 101 and 102.
[321] Ofcom analysis based on [✂]; and [✂]. Based on [✂] data, we estimate that around [✂]% of UK customers ([✂] customers) that renegotiated their deal received a [✂], and around [✂]% of UK customers ([✂] customers) that renegotiated received the [✂]. For [✂]. For the customers that we were able to observe this for ([✂] customers), around [✂]% ([✂] customers) received a lower discount upon renegotiation, despite increasing their spend. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂].
[322] [✂].
[323] [✂].
[324] Context Consulting research report, slide 100.
[325] [✂].
[326] [✂].

deduction on the discounts received by customers; and customers potentially being subject to egress fees when moving data out of a given cloud provider.

A5.15 In our view, it is not feasible to use the data supplied by the hyperscalers to conclusively identify whether or not barriers to switching and multi-cloud lead to worse customer outcomes. This is discussed further in the section below.

Limitations of the regression analysis

A5.16 The hyperscalers, [✂] and [✂], submitted a regression analysis to indicate that on average, committed spend discount customers do not appear to receive worse discounts upon renegotiations, [✂]. Based on this, they argue that [✂].

A5.17 However, we think that the evidence submitted by the hyperscalers is not sufficiently comprehensive to allow us to test whether customers that face greater barriers to switching or multi-cloud are being harmed. As such, we place limited weight on these findings. The limitations include issues related to the scope of the data, as well as a range of methodological issues.

Limitations on the scope of the data

A5.18 There are limitations associated with the scope of the data which mean we cannot fully test whether customers are being harmed due to a limited ability to switch or multi-cloud. The data submitted by the hyperscalers only captures the committed spend discounts negotiated by customers and does not capture other elements of pricing relevant to committed spend discount customers. Based on the datasets submitted by [✂] and [✂], we therefore do not think that we have the full data needed to assess whether customers who face greater barriers receive worse outcomes upon renegotiations.

A5.19 We are aware that committed spend discount customers may be able to negotiate other forms of discounts (e.g. [✂]),[327] alongside their committed spend discount. For example, one hyperscaler, [✂], explained that customers renegotiating at the same commitment spend may sometimes receive [✂] discounts due to the specific circumstances of the deal, and [✂]. The hyperscaler, [✂], estimates that incentives such as [✂] are provided to their committed spend customers in less than [✂]% of their deals.[328] This implies there is also scope for worse outcomes via these other sources of discounts which we cannot observe for in the data, such that two customers with a similar level of commitment, may be subject to different additional discounts. Additionally, we are also aware that the hyperscalers offer egress discounts to a small number of customers (see Section 5, subsection on data egress fees). We think that there is likely to be scope for customers to negotiate on these additional discounts, which means that observing the committed spend discounts of customers alone may not enable us to fully assess the outcomes customers may face upon renegotiations.

A5.20 In addition, the realised committed spend discount for these customers could be lower than the agreed discounts observed in the datasets. For example:

i) We understand that some committed spend discount customers are required to take up ancillary services, such as [✂] which results in an effective reduction on

[327] [✂] response dated [✂] to our follow-up email dated [✂], question [✂], concerning the s.174 notice dated [✂]; and [✂] response dated [✂] to our follow-up email dated [✂], question [✂], concerning the s.174 notice dated [✂].
[328] [✂].

their committed spend discount. We consider this further below in our customer engagement subsection.

ii) The discounts observed in the dataset are agreed discounts and may not actually be the discount realised if customers are unable to meet their commitment (i.e. their actual spend is lower than their committed spend). The hyperscalers recognise that a small proportion of customers tend to miss their commitment.[329] Based on the completed contracts for which we could observe actual spend, we estimate that around [✂]% of [✂] customers ([✂] customers) and [✂]% of [✂] customers ([✂] customers) may have missed their commitment over their contract period, with most missing it by more than [✂]%.[330] However, these estimates should be treated as an upper bound, as we understand that the data on actual spend provided by [✂] and [✂] may understate their customers spend.[331]

Methodological limitations

A5.21 There are also several methodological limitations to the regression analyses presented by [✂] and [✂], which means they cannot identify whether customers that face greater barriers to switching or multi-cloud are being harmed. This is for a number of reasons.

A5.22 First, the renegotiation variable is not able to distinguish between customers that do and do not face barriers to switching and multi-cloud. As a result, it cannot reliably identify whether there are customers that face significant barriers to switching and multi-cloud, and whether this results in worse outcomes for those customers. This is for several reasons:

i) It is plausible that some customers may already face barriers to switching or multi-cloud to some extent before they sign their first committed spend discount contract. We understand that some customers are likely to have taken up PAYG prices, before reaching the volumes required for a committed spend discount.[332] Even as PAYG customers, many are likely to already face some barriers to switching and multi-cloud. Therefore, examining the discounts received by customers when negotiating a committed spend contract for the first time, and comparing this to the discounts received by customers when renegotiating their contracts, may not accurately reflect the difference between a customer that does and one that does not face barriers to switching or multi-cloud.

ii) If some customers do not renew their contract due to having weaker bargaining power, then they will disappear from the data sample. In this case, the difference in discounts observed between the first and renewed contracts will be subject to selection bias: we would observe discounts for all customers on their first contract, but only for a subset of customers for their second and subsequent contracts.[333]

---

[329] [✂] estimate that out of their customers with completed contracts in 2022, [✂]% of their committed spend customers were able to meet at least [✂]% of their commitment during their contract term. [✂]. [✂] estimate that less than [✂]% of their committed spend customers were unable to meet their commitment. However, the data time frame for this is unclear. [✂].

[330] Ofcom analysis based on [✂]; and [✂].

[331] We understand from [✂] and [✂] that their actual spend may take into account the committed spend discount of customers, or potentially other discounts offered to customers. The actual spend of these customers before the discounts are applied could be higher and therefore, some may have had the revenues to meet or exceed their commitment spend. [✂] response dated [✂] to our follow-up email dated [✂], question [✂], concerning the s.174 notice dated [✂]; and [✂] response dated [✂] to our follow-up email dated [✂], question [✂], concerning the s.174 notice dated [✂].

[332] Ofcom / [✂] meeting, [✂]. Ofcom / [✂] meeting, [✂].

[333] For example, we observe that [✂]. Ofcom analysis based on [✂].

iii) There are limitations associated with identifying renegotiating customers in the datasets, which further makes it difficult to test for harm using renegotiation as a proxy for the barriers to switching and multi-cloud customers face. Amongst committed spend discount customers in the UK, the hyperscalers are unable to accurately identify the first committed spend contract for all customers, and therefore accurately distinguish between new and renegotiating customers as part of their analysis. For example, one hyperscaler, [✂], suggest that some customers identified to have taken a contract for the first time may actually be renegotiating customers, although we do not have an estimate for the proportion of customers that may be mis-classified in this way.[334] Similarly, we are also aware that some first-time committed spend contracts in [✂] dataset may not actually be the first-time contract of a customer. For [✂], we estimate that this issue may affect up to [✂]% of their deals, and [✂]% of their customers in the dataset.[335]

A5.23 Second, the regression analyses submitted by the hyperscalers may be subject to endogeneity[336] issues which could further impact the robustness of their regressions. This is because the regression estimates the effect of various variables on the discount level of customers, such as commitment spend, contract duration and customer renegotiation. However, customers are likely to negotiate on discounts, duration, and commitment spend all simultaneously. The simultaneity between the dependent variable (i.e. discounts) and independent variables (i.e. committed spend, duration) may bias the renegotiation coefficient, meaning we cannot attach weight to the sign and significance of the renegotiation coefficient.

A5.24 Third, the renegotiation variable may not account for the total effect of bargaining power on discounts. This is because the bargaining power of customers can have an influence on both discounts directly, as well as indirectly (through contract duration and commitment spend). For instance, a customer's bargaining power may have an influence on discounts, such that a customer with stronger bargaining power may be able to negotiate a higher discount. At the same time, a customer's bargaining power may also influence their commitment spend and contract duration, which in turn also determine their discount level. The hyperscalers' regressions, which [✂], may therefore not be able to attribute the potential indirect effects on discounts to bargaining power. Therefore, the coefficient on the renegotiation variable may be biased towards zero, as the total effect of bargaining power on discounts might be different from the partial effect that is captured by the regressions.

A5.25 Fourth, we are also aware that the data presented by the hyperscalers has limitations and in particular, that there are issues with the quality of [✂] data which makes it less robust. The hyperscaler's dataset only contains a small number of UK customers that have renegotiated their committed spend discount (around [✂] customers). In addition, they suggest that their dataset is incomplete and could be missing at most [✂]% of their UK committed spend

[334] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[335] Ofcom analysis based on [✂].
[336] Endogeneity issues arise in regression models, when the effect of the independent variable (i.e. committed spend) on the dependent variable (i.e. discount) may be subject to biases, which means causal inferences between the two variables cannot be drawn. In the hyperscalers' regressions, the simultaneity of the dependent and independent variables, in other words the ability for both variables to impact one another, may be a cause of endogeneity.

discount customers.[337] Therefore, their regression analysis at the UK level may not necessarily be comprehensive due to the missing data, or robust due to the small number of observations. The hyperscaler themselves recognise the small number of observations as a weakness to their regression analysis for the UK. Although, their analysis at the global level is based on a larger number of observations, like the UK dataset, it may also be incomplete,[338] which is likely to be a weakness to their regression analysis at the global level.

A5.26 Finally, even when controlling for commitment spend and duration as part of the regressions for the UK, we observe that the R-squared for [✂], which suggests that there remains a lot of unexplained variation (more so for [✂]) in determining discount levels.[339] When [✂] control for "customer-specific characteristics" in one of their regression specifications, we observe a [✂] in the R-squared.[340] This suggests that there are likely to be other factors (beyond spend and duration) that are important in determining the level of discounts received by customers, and their regressions do not account for this.

A5.27 Overall, taking into account the combination of limitations associated with the regressions, we place limited weight on this analysis. Given that it is beyond the scope of the market study to evidence such harm and the difficulties in testing conclusively whether customers who face barriers to switching and multi-cloud are harmed, we consider that it is not for the market study to resolve these issues or advance the analysis that has been presented.

Variation in discounts received by customers

A5.28 As discussed above, we are aware that even when accounting for spend and contract duration, there appears to be a lot of unexplained variation in the discounts received by customers. This suggests that other factors besides spend and contract duration may also be important in determining the discounts received by customers. In fact, the evidence received from the hyperscalers suggests that there is scope for customers to negotiate on discounts, which does suggest that other factors may also come into play when determining discounts. One of these factors could be the degree of bargaining power the customer holds and hence, the variation we observe in the discounts received by customers could be the outcome of differences in the bargaining power of customers. This means that customers who have weaker bargaining power due to having a limited ability to switch or multi-cloud could be subject to lower discounts at a given spend level, relative to a customer with stronger bargaining power.

A5.29 Below, we consider the evidence on the extent of variation in the discounts received by committed spend customers and also other groups of customers that receive discounts, such as [✂] customers. We also consider evidence on the scope that these customers have to negotiate more individualised discounts.

[337] [✂] response dated [✂] to our follow-up email dated [✂], question [✂], concerning the s.174 notice dated [✂].
[338] [✂].
[339] [✂]; and [✂].
[340] [✂].

Variation in the discounts received by committed spend customers

A5.30 We acknowledge that customers with privately negotiated committed spend discounts are typically able to obtain a higher percentage discount if they increase the level of spend and the duration of their commitment. Based on [✂] and [✂] analysis, this appears to be the general trend. For example, the data suggests that around [✂]% of [✂] renewals and [✂]% of [✂] renewals resulted in committed spend discount customers spending more with each contract and receiving a larger discount in return.[341] One hyperscaler further suggested that they have a standard discount schedule in place based on the level of spend and duration a customer commits to.[342]

A5.31 However, amongst committed spend discount customers, we also observe that customers with similar levels of spend and contract duration can receive different discounts. Figure A5.1 below illustrates the wide distribution of discounts received by committed spend discount customers with a contract duration of [✂] months. We specifically examine discounts with these contract durations, as these are the contract lengths commonly taken up by customers.[343] For example, one hyperscaler's, ([✂]), data suggests that customers with a commitment spend of [✂] million and a contract duration of around [✂] months can receive discounts ranging from around [✂], while [✂] data suggests that customers with a commitment spend of [✂] million and a contract duration of around [✂] months can receive discounts ranging from around [✂].[344]

Figure A5.1: [✂] committed spend and discounts for contracts with a [✂] months duration (left) and [✂] committed spend and discounts for contracts with a [✂months duration (right)

[✂] [✂]

Source: Ofcom analysis based on [✂]; and [✂].

A5.32 The variation we observe suggests that there may be other factors, besides spend and duration, determining the discounts received by committed spend discount customers. The hyperscalers themselves acknowledge that there is scope for these customers to negotiate more individualised discounts, which suggests that they do deviate from their standard discount schedules. For example, [✂] suggest that committed spend discounts tend to be "heavily negotiated with customers, [✂]. They also mention that [✂].[345] More generally across their customers, [✂] also indicate that customers tend to negotiate on a range of terms, including on discounts. They further suggest that for the majority of their cloud revenues, it is customer negotiation that drives the end prices and discounts received by customers.[346]

---

[341] Ofcom analysis based on [✂]; and [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂].
[342] [✂].
[343] However, we also observe variation in the discounts received by committed spend discount customers at other contract lengths too.
[344] Ofcom analysis based on [✂]; and [✂]. [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; and [✂].
[345] [✂].
[346] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

A5.33 Therefore, the evidence above suggests that there is wide variation in the discounts received by committed spend discount customers who make similar commitments. We think this variation may be reflective of differences in the bargaining power customers can exercise as part of their negotiations with the hyperscalers. This suggests that there may be scope for worse outcomes amongst those customers who face greater barriers to switching and multi-cloud, and therefore have weaker bargaining power.

Variation in the discounts received by other groups of customers

A5.34 Besides committed spend discounts, customers are also able to receive other types of discounts. In response to the interim report, one hyperscaler, [✂], submitted information on the pricing outcomes of their [✂] customers.[347] These customers account for around [✂]% of their customer base and [✂]% of their cloud revenues in the UK.[348] They identify these customers as a distinct category, separate from PAYG customers that are subject to website prices and customers with privately negotiated committed spend discounts.

A5.35 The hyperscaler, [✂], suggest that its discounting policy for [✂] customers is "standardised", such that customers tend to qualify for a given maximum discount as set out in their guidelines by committing to a specific consumption tier, and almost all customers receive the same discounts before and after contract renewal.[349] [✂].[350] [✂].

A5.36 The hyperscaler, [✂], suggest that their committed spend discount programme is the primary way for customers to receive individualised discounts. We note that [✂] does offer lower pricing to customers through its reserved instances, spot instances and savings plans, but it suggests that its discounting is automatically generated for these customers based on their spend levels.[351]

A5.37 Therefore, the data submitted by [✂] suggests [✂]. With [✂], the indication is that the committed spend discount programme is the primary way for customers to receive individualised discounts and hence there is little scope for other customers to negotiate more individualised discounts, but we have not tested this further with evidence or data.

Customer engagement

A5.38 Through our engagement with customers, we have heard a variety of examples that demonstrate the inflexibility some customers face in negotiations, as well as the variation in outcomes that customers experience. Customers often feel this is a result of the imbalance in bargaining power between themselves and the hyperscalers; and in some cases directly attribute this to their reliance on a single provider and inability to switch. This can lead to different customers paying different prices for similar purchases, but we have also heard other concerns related to bargaining power imbalances. This may present itself as an inability to negotiate on certain contractual terms, mandatory take-up of ancillary services, or pressure to increase spending commitments, with a resultant impact on the effective

[347] [✂].
[348] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].
[349] [✂].
[350] [✂].
[351] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂].

prices customers pay – these are outlined in turn below. These experiences are not universal, which suggests that the relative bargaining power between provider and customer is important when negotiating contracts.

A5.39 We have heard that there are limitations on what customers are able to negotiate with providers, which they have in some cases associated with a weak bargaining position. As outlined in section 4, we have heard examples of customers being unable to negotiate on fixed prices for the duration of their contracts,[352] or make any changes to standard terms and conditions that are generally negotiable in other markets.[353] For example, the one-way protection (in the provider's favour) from exchange rate fluctuations resulted in prices changing by up to 10% in recent bills, and was inconsistent with other supplier relationships.[354] This experience of limited negotiating power is not restricted to smaller companies. We have heard from large companies that they do not have sufficient bargaining power to negotiate fully on terms such as prices and discounts with their cloud provider.[355] Some customers associated their relatively weak bargaining position with their increasing dependence on a single cloud provider and limited ability to switch providers.[356]

A5.40 For some [✂] customers with a cross-service committed spend discount, the automatic take up on ancillary services [✂] has been another issue some customers have struggled to negotiate on. This mandatory take-up of ancillary services has been reported in some cases to wipe away the majority of the benefit from a hard-fought discount.[357] The hyperscaler [✂] confirmed that customers with committed spend discounts [✂], however they noted that customers that do not want [✂].[358]

A5.41 As discussed in further detail in Section 5, some customers also told us that they had faced pressure from one provider to commit to greater spend levels, in some instances to just maintain their existing discounts. In other instances, customers experienced pressure on their rate of growth[359] and some customers noted having concerns around having to revert to paying list prices in the future if they failed to continue to grow their spend when renewing their contract.[360] For some customers, agreeing to increased spend commitments could lead to an over-commitment, where their final usage and spend is less than their minimum committed spend. Customers also raised a general concern around a lack of

---

[352] [✂] confidential response to the interim report, page [✂].
[353] [✂] confidential response to the interim report, page [✂], BT Group response to the CFI, page 12. Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[354] Microsoft website. Consistent global pricing for the Microsoft Cloud [accessed 31 August 2023]; UC Today, 2023. Costs to Increase as Microsoft Sets Consistent Global Pricing for Cloud Services [accessed 14 September 2023].
[355] [✂] response dated [✂] to s.174 request dated [✂]; [✂] response dated [✂] to s.174 request dated [✂], question [✂]; and Ofcom / [✂] meeting, [✂], subsequently confirmed by email by [✂] on [✂].
[356] [✂] response to the call for inputs, page [✂]; and [✂] response dated [✂] to s.174 request dated [✂], question [✂].
[357] One customer [✂] estimated that being required to purchase [✂] could represent a deduction of up to 60% on their discount. Similarly, another customer [✂] also indicated that a large proportion of their committed spend discount was cancelled out by the non-negotiable take up of [✂]. Email from [✂] received [✂] subsequently confirmed at Ofcom / [✂] meeting, [✂] and Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email dated [✂].
[358] [✂].
[359] Ofcom / [✂] meeting, [✂], subsequently confirmed by email by [✂] on [✂]; Ofcom / [✂] meeting, [✂], subsequently confirmed by email on [✂]; and Ofcom / [✂] meeting, [✂], subsequently confirmed by email on [✂].
[360] Ofcom / [✂] meeting, [✂] subsequently confirmed by [✂] by email on [✂].

transparency in the negotiation process, as they were unclear how a change in size or duration of their commitment would affect the size of their discount. Nearly all customers we spoke to told us that it is common for the cloud provider to present their own committed spend forecasts for customers, which customers then have to work to negotiate down. In contrast, another customer told us that they felt the same provider would not allow them to over-commit on their spend, as they felt their provider was unlikely to risk their long-term relationship.[361] One customer acknowledged that there tends to be scope to negotiate on discounts, as well as the ability to receive credits from hyperscalers, but these depend on the leverage that customers may have.[362]

A5.42 Despite these challenges, some customers told us that they have taken steps to strengthen their ability to negotiate with cloud providers. We heard that one customer had undertaken a rearchitecting of their entire platform, so it could be effectively delivered through a multi-cloud deployment. In doing this, the customer has reduced their spend with their main cloud provider and created a credible threat to reduce spend further over time. In turn, they have been able to begin negotiations on terms that were initially and in previous contracts deemed non-negotiable.[363] Whilst this enabled the customer to increase their bargaining position, this process has taken significant time and cost. We also note that for this customer rearchitecting was incidental to other work needed in the organisation. We recognise that feedback we have received from the majority of customers suggests that this approach is not feasible for most.[364]

A5.43 The evidence we have received demonstrates the spectrum of bargaining ability and outcomes for customers. We have heard from some customers that they are unable to negotiate favourable direct and indirect price terms with providers, due to the bargaining power imbalance they face. Some customers appear to experience complete inflexibility in negotiating with providers, whilst others are able to reach some compromise. While some have managed to strengthen their negotiating position, this experience is limited and challenging to achieve.

[361] Ofcom / [✂] meeting, [✂], subsequently confirmed by [✂] by email on [✂].
[362] Ofcom / [✂] meeting, [✂] subsequently confirmed by email by [✂] on [✂].
[363] Ofcom / [✂] meeting, [✂] subsequently confirmed by email by [✂] on [✂].
[364] For example, [✂] said that IAM had been a significant barrier in taking up Google in addition to AWS, as they found it technically challenging and potentially unsafe to manage both AWS and Google IAM due to the high degree of differentiation. [✂] said that avoiding the complexities associated with managing ancillary cloud services (e.g. IAM or billing) across clouds was one of the key factors driving its decision to concentrate its cloud usage on AWS.

A6. Terms of reference

The Office of Communications (Ofcom) in the exercise of its powers under sections 131 and 133 of the Enterprise Act 2002 (EA02) as provided for by section 370(3A)(b) of the Communications Act 2003 (CA03) read together with section 130A EA02 hereby makes a reference to the Chair of the Competition and Markets Authority (CMA) for the constitution of a group under Schedule 4 to the Enterprise and Regulatory Reform Act 2013 for an investigation in relation to public cloud infrastructure services in the United Kingdom.

Ofcom has reasonable grounds to suspect that a feature or a combination of features of the markets for the supply of those goods and services in the United Kingdom prevents, restricts or distorts competition. In particular, conduct which may create barriers to switching and multi-cloud.

For the purposes of this reference:

- 'Cloud infrastructure services' means services that provide access to processing, storage, networking, and other raw computing resources (often referred to as infrastructure as a service, IaaS) as well as services that can be used to develop, test, run and manage applications in the cloud (often referred to as platform as a service, PaaS).
- 'Public cloud computing' means a cloud deployment model where cloud services are open to all customers willing to pay, and computing resources are shared between them.
- 'Multi-Cloud' means a cloud deployment model involving the use of more than one public cloud provider by a single customer, where multiple clouds may or may not be integrated with each other.

A7. Glossary

Term	Definition
API (application programming interface)	A software interface that allows two or more pieces of software to communicate with each other.
Application portability	Ability to migrate an application from one cloud to another or between a customer's IT environment and a cloud and be able to run it correctly in the target cloud with minimal disruption.
AWS (Amazon Web Services)	A subsidiary of Amazon Inc that provides a full range of cloud services at scale to UK customers.
Bare metal services	Services which offer access to dedicated servers with no or limited software installed (e.g. no operating system or virtualisation).
CFI (Call for inputs)	Ofcom's publication entitled 'Cloud services market study - Call for inputs' dated 6 October 2023.
Cloud	A cloud is a suite of cloud computing services (IaaS, PaaS and/or SaaS) hosted in data centres, provided by a given cloud provider (e.g. AWS, Azure or Google Cloud platform). Note, cloud is sometimes used to refer to cloud computing.
Cloud computing	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.
Cloud ecosystem	A portfolio of services across the service and deployment models of a cloud, including its marketplace.
Cloud infrastructure services	Services that provide access to processing, storage, networking, and other raw computing resources (often referred to as infrastructure as a service, IaaS) as well as services that can be used to develop, test, run and manage applications in the cloud (often referred to as platform as a service, PaaS).
Cloud marketplace	A website operated by a cloud provider where customers can purchase services (supplied by the marketplace owner and third parties) that are compatible with their clouds.
Cloud provider	Cloud providers are vertically integrated suppliers of cloud services that operate their own cloud infrastructure (i.e. they own the underlying raw computing resources).
Cloud services	All services involved in the provision of cloud computing.

Computing resources	Physical or virtual components within an information technology system, including servers, storage, network and applications.
Container as a service (CaaS)	A layer where applications or parts of applications run separately in a container, but sections of the operating system and storage are shared.
Container	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).
Content delivery network (CDN)	A CDN is a geographically distributed network of servers aiming at fast delivery of internet content, including HTML pages, scripts, stylesheets, images, audio files and videos. Serving content from a CDN, located closer to the end-user, may improve load times, reduce bandwidth costs and increase availability.
Data centre	Buildings that house hardware needed for cloud computing such as servers and network equipment.
Data portability	Ability to easily move data from one cloud to another or from a customer's IT environment to a cloud and have that data usable in the target cloud with minimal disruption.
Database as a service (DBaaS)	A cloud service that provides customers with access to a database.
Data sovereignty	Refers to the concept that data is subject to the laws and regulation of the country in which data are collected, processed and stored.
Disaster recovery as a service (DRaaS)	A cloud computing service model that allows an organisation to back up its data and applications (for example, in another region served by the cloud provider), and provide disaster recovery orchestration through a SaaS solution.
Duplicated multi-cloud (cloud duplication)	This multi-cloud architecture 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.
Edge cloud or multi-access edge computing (MEC)	Processing workloads and storing data close to the edge of a telecoms network, i.e., the physical location where users connect with the telecoms network.
Egress fees	Fees charged by cloud providers to customers for data leaving the cloud network in transit to an external location.
GB	Gigabyte.
GBP (£)	British pound sterling.

Google	A subsidiary of Alphabet Inc that provides a full range of cloud services (Google Cloud Platform) at scale to UK customers.
Hybrid cloud computing	A cloud deployment model involving a combination of public clouds and private environments (such as private clouds or on-premises resources). which allow workloads to be shared between them.
Hyperscalers	AWS, Microsoft and Google.
IaaS (infrastructure as a service)	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.
Independent software vendor (ISV)	Supplier of cloud services, typically PaaS and/or SaaS, that does not own any of the underlying raw computing resources.
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. In doing so, the customer integrates different customer applications, customer data, and/or cloud services hosted on two or more public clouds into a consolidated architecture.
Interim report	Ofcom's publication entitled 'Cloud services market study. Interim report' dated 5 April 2023.
Internet of things (IoT)	The network of devices that contain the hardware (including sensors and actuators), software and firmware which allow the devices to connect, interact, and freely exchange data and information.
Interoperability	The ability of computer systems or software to communicate with one another.
kW	Kilowatt.
Market study notice	Ofcom's published legal notice of 6 October 2022, pursuant to section 1340A of the Enterprise Act 2022 as amended and applied by section 370 of the Communications Act 2003, launching a market study into the provision of cloud services in the UK.
Microsoft	Microsoft Corporation, a company that provides a full range of cloud services (Azure) at scale to UK customers.

MIR consultation	Ofcom's publication entitled 'Public cloud infrastructure services. Consultation: Proposal to make a market investigation reference' dated 5 April 2023.
MIR notice	Ofcom's published legal notice of 5 April 2023, pursuant to section 131B(1) of the Enterprise Act 2002, proposing a market investigation reference under section 131 of the Enterprise Act 2002.
Multi-cloud	A cloud deployment model involving the use of more than one cloud by a single customer, where multiple clouds may or may not be integrated with each other.
MW	Megawatt.
On-premises IT	Refers to IT infrastructure (hardware and software) that is hosted on the premises of the person or organisation using the software, rather than at a remote facility. On-premises IT could be part of a traditional IT or a cloud architecture.
Open-source software (OSS)	Software released under a license in which the copyright holder grants users the right to freely use, change, and distribute the software and its source code.
PaaS (platform as a service)	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. The virtual environment, the underlying infrastructure and computing resources are typically owned and managed by the same service provider and are typically hidden from the consumer.
Private cloud computing	A cloud deployment model where computing resources are dedicated to (as opposed to shared between) individual customers.
Public cloud computing	A cloud deployment model where cloud services are open to all customers willing to pay, and computing resources are shared between them.
ROCE (return on capital employed)	Operating profit (measured as earnings before interest and tax) divided by capital employed.
SaaS (software as a service)	Complete applications hosted in the cloud. They can be offered by the cloud provider that owns the underlying cloud infrastructure or by an independent software vendor.
Siloed multi-cloud	This multi-cloud architecture 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).
Stack	A set of hardware and software components that work together to create a computing platform for running applications.
Switching between clouds	Switching services from one public cloud to another.
Switching within cloud	Switching between cloud services hosted on the same public cloud.
TB	Terabyte.
Traditional IT	Dedicated physical computing resources that are not part of the cloud. These are typically owned by, and located on the premises of, the customer.
Usage commitments	Where customers receive cheaper rates compared to on-demand for committing to a certain configuration or amount of usage for a specified term, typically 1-3 years.
USD ($)	United States dollar.
Vertically integrated providers	All cloud providers, i.e., the hyperscalers and smaller cloud providers.
Virtualisation	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 unit, random access memory and storage) to be divided into multiple virtual servers, commonly called virtual machines.
Virtual machines (VMs)	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.
WACC (weighted average cost of capital)	The rate that a company is expected to pay on average to all its security holders, both debt and equity, to finance its assets.
Workload	A specific application, service, capability or a specific amount of work that can be run on a cloud resource.