Cloud services market study

Interim report, annexes 5 to 8

Non-confidential version – redacted for publication [✂]

CONSULTATION:

Publication date: 5 April 2023

Closing date for responses: 17 May 2023

Contents

Annex

A5. Shares of supply for cloud infrastructure services in the UK

A5.1 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).

A5.2 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]

A5.3 We also comment on which types of products represent most revenue for the hyperscalers in the UK.

Source of data

A5.4 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').

A5.5 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

A5.6 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.

A5.7 We asked these 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]

A5.8 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).

A5.9 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 IaaS/PaaS categorisation. Similarly, while all companies provided estimates of UK revenues,

[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.

in some cases these did not align to what customers were billed or how revenue was recognised in statutory accounts.[3]

A5.10 The table below summarises the revenue data provided by the 11 companies we sent information requests to.

Table A5.1: 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.

A5.11 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.

Table A5.2: 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 public 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.

[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.

Limitation	Adjustments to the data
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, [✂].[4] Microsoft and Google were not able to quantify the amount of compute and storage revenues associated with usage by these services.[5] 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. In particular, as we understand that Microsoft records all IaaS fees against IaaS services and AWS and Google do not, IaaS revenues for AWS and Google could appear relatively low, and their PaaS revenues relatively high compared to Microsoft. For this reason, UK revenues and shares of supply for IaaS and PaaS combined may be less affected by differences in choices on how to record revenues.
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.

[4] AWS response dated 17 February 2023 to question 2 of our follow up email dated 28 January 2023 concerning the s.174 notice dated 24 October 2022.
[5] 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 question 2 of our follow up email dated 27 January 2023 concerning the s.174 notice dated 26 October 2022.

Limitation	Adjustments to the data
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.

Synergy and IDC UK data

A5.12 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.

Synergy

A5.13 We bought UK revenue data on IaaS and PaaS from Synergy for the calendar years 2018 to 2021.[6]

A5.14 Synergy's IaaS data is broken down into storage, compute, networking, and other.

A5.15 Synergy's PaaS data is broken down into analytics, databases, IoT (internet of things), and other.

A5.16 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

A5.17 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).[7] This data covered the calendar years 2017 to 2021.

[6] Synergy Research Group, 4Q 2021 Cloud Infrastructure Services United Kingdom Market Share Report, March 2022. 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.
[7] IDC, Public Cloud Services Tracker 2021 H2 (published April 2022). We converted IDC's data from US dollars to pound sterling using average monthly exchange rates from ONS for the relevant period.

A5.18 IDC's IaaS data is broken down into three secondary markets: storage, compute, and networking.

A5.19 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.

A5.20 We did not purchase any additional granularity on the secondary markets for SaaS - Systems Infrastructure Software or SaaS - Applications.

A5.21 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.[8]

UK revenues for IaaS and PaaS

Methodology

A5.22 To estimate total UK revenues for IaaS and PaaS we did the following:

- 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 A5.2 above. Total UK revenue from these 11 companies in 2021, across both IaaS and PaaS, was around £4.0bn.
- 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).

A5.23 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.

A5.24 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:

- Potential differences in definitions of IaaS, PaaS and SaaS, and views into which category individual services fit,[9]
- differences in the number of companies tracked, and
- potential differences in source data, modelling methodology, and approach to estimating UK revenues.

[8] IDC Worldwide Software Taxonomy, 2022 and IDC Worldwide IT Cloud Services Taxonomy, 2019.

[9] 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.

A5.25 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, [✂].[10] For publication, we report these averages within ranges of £500m.

UK revenues for IaaS and PaaS

A5.26 We estimate that in 2021, cloud infrastructure services generated revenues of £4.5bn to £5.0bn. Between 2019 and 2021, UK revenues for IaaS and PaaS combined grew by 25% to 30% per year.

Table A5.3: UK IaaS and PaaS revenues, 2019 – 2021 (£bn)

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

Source: Ofcom analysis of data provided in response to our information requests and data from Synergy and IDC. Annual growth based on the compound annual growth rate between 2019 and 2021.

A5.27 Within IaaS, the main product categories are compute, storage and networking. In 2021 we estimate that these represented around 65%, 20% and 8% of UK IaaS revenues respectively.[11] 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.

A5.28 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.

UK shares of supply for IaaS and PaaS

Methodology

A5.29 To estimate UK shares of supply for IaaS and PaaS we did the following:

- 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

[10] [✂].

[11] 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.

we also comment on the shares of supply for some of the larger providers outside of the hyperscalers.

- 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 A5.3.
- 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%).

UK shares of supply for IaaS and PaaS

IaaS

A5.30 The table below shows our estimated shares of supply for UK IaaS.

Table A5.4: UK IaaS shares of supply, 2019 – 2021

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

A5.31 We estimate that AWS and Microsoft represent approximately [✂] [70% to 80%] of UK IaaS revenues, a percentage that has remained broadly stable since 2019. Within this, we estimate that Microsoft's share of UK IaaS revenues has grown slightly while AWS's share has reduced slightly. As noted in Table A5.2, 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.

A5.32 While Google's IaaS revenues in the UK have grown quickly since 2019, in 2021 we estimate it represented [✂] [0% to 5%] of UK IaaS revenues, significantly behind AWS and Microsoft.

A5.33 Within the 'other' category, IBM and Oracle appear to have the next largest shares of supply, though these are both [✂] [0% to 5%]. Over this period, UK IaaS revenue growth for smaller cloud providers broadly kept up with hyperscaler revenue growth, although from a lower base.

PaaS

A5.34 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 2021, only the hyperscalers and IBM were active in each of IDC's seven 'secondary markets'[12] while Oracle operated across most of these categories. The majority of other companies tracked by IDC were only active in one or two secondary markets.

A5.35 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, looking at PaaS overall may mask that there could be more alternatives for some types of service than others.

A5.36 The table below shows our estimated shares of supply for UK PaaS.

Table A5.5: UK PaaS shares of supply, 2019 – 2021

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

A5.37 We estimate that AWS and Microsoft represented [✂] [40% to 50%] of UK PaaS revenues in 2021 – a lower share than for IaaS. Within this, we estimate that Microsoft's share of UK PaaS revenues has grown slightly while AWS's share has remained stable. As noted in Table A5.2, 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.

A5.38 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 2021 and its share increasing since 2019. [✂].

A5.39 The "other" category is very broad and includes many companies with low shares of supply. Oracle and MongoDB have some of the larger shares of supply of companies in the 'other' category, followed by IBM; we estimate they represented around [✂] [0% to 5%], [✂] [0% to 5%] and [✂] [0% to 5%] respectively of UK PaaS revenues in 2021. Over this period, UK PaaS revenue growth for smaller cloud providers generally lagged behind hyperscaler revenue growth.

[12] 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.

A5.40 We estimate that ISVs collectively accounted for [✂] [30% to 40%] of UK PaaS revenues in 2021.

IaaS and PaaS combined

A5.41 Table A5.6 shows our estimated shares of supply for UK IaaS and PaaS combined, drawing on the information presented above. We estimate that in 2021 AWS and Microsoft between them had [✂] [60% to 70%] 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 2021.

A5.42 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 2021.

A5.43 Over this period, we estimate that AWS's UK share of supply fell marginally, Microsoft's grew slightly and Google's share experienced stronger growth, although from a lower revenue base.

Table A5.6: UK shares of supply for IaaS and PaaS combined. 2019 – 2021

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

Major product categories for the hyperscalers

A5.44 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.

A5.45 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.[13] Based on the information provided, we identified revenues

[13] Microsoft noted that some services could reasonably fit into multiple product categories. Microsoft response dated 9 December 2022 to question 1 of Part B of the s.174 notice dated 21 October 2022.

associated with five broad categories: compute, storage, networking, data management and analytics.

A5.46 Figure A5.7 shows the proportion of UK IaaS and PaaS revenues that we estimate was associated with these five categories in 2021.[14] It indicates that:

- Compute is typically the largest product category
- Data management services are typically the second or third largest product category
- Storage is relatively important for Microsoft
- [✂] for GCP

Figure A5.7: Proportion of hyperscalers' UK IaaS and PaaS revenue from particular product categories, 2021

[✂]

Source: Ofcom analysis of data provided in response to our information requests.[15] Note that the chart does not include an 'other' category so does not sum to 100%.

A5.47 We estimate that compute was the slowest growing category in the UK across the hyperscalers between 2019 and 2021 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.

A5.48 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.

[14] 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.

[15] Our analysis was based on the following responses: AWS response dated 9 December 2022 to question 4 of Part B of the s.174 notice dated 24 October 2022; Google response dated 16 December 2022 to questions 1 and 4 of Part B of the s.174 notice dated 26 October 2022; Microsoft response dated 9 December 2022 to questions 1 and 4 of Part B of the s.174 notice dated 21 October 2022.

A6. Profitability of the hyperscalers

Introduction

A6.1 In this annex we set out 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.

A6.2 Profitability can be one indicator of how well competition is working. We set out our analysis of the following in this annex:

- 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.
- We assess whether the hyperscalers' cloud businesses have generated returns persistently above their cost of capital. Evidence of this may indicate limitations to the competitive process, including the degree to which larger firms face competitive constraints, and potential barriers to entry or expansion by competitors. 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.

A6.3 This annex is split into the following sections:

- **Source of data and scope of analysis**
- **Hyperscaler cloud profitability**: We assess hyperscalers' cloud revenues and operating profits and compare these to other cloud infrastructure providers.
- **Return on capital employed versus cost of capital**: We assess the return on capital employed (ROCE) earned by AWS and Microsoft on their cloud investments and compare this against the cost of capital.

Source of data and scope of analysis

Source of data

A6.4 We have run our analysis at a global level because the major cloud providers are globalised businesses, with many of their expenses and investments in cloud services serving their global customer base. 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 for 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.[16]

A6.5 Our analysis is based on global information from published financial statements and information provided by cloud providers in response to our statutory information requests.

A6.6 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.

A6.7 Table A6.1 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.

[16] In recent market studies, the CMA has also considered global businesses when assessing profitability. 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 A6.1: 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[17] – Intelligent Cloud also includes non-cloud server products and enterprise support and consulting services
IBM	IBM public cloud[18]	N/A (IBM reports revenue for 'hybrid cloud', which appears to be included in 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-premise products like Oracle Database[19]
OVHcloud	Public Cloud	Public Cloud	No
Alibaba[20]	Alibaba Cloud	Cloud	Yes – some activities within the reported Cloud segment do not relate to infrastructure services[21]
DigitalOcean	DigitalOcean (whole business)	DigitalOcean (whole business)	No

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

A6.8 Table A6.1 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 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 specifically.

---

[17] 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).

[18] 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.

[19] 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.

[20] Alibaba Cloud often appears as one of the largest cloud providers globally given its position in China. For example, IDC ranked it third globally in the provision of Foundational Cloud Services (which includes the IaaS, PaaS and SaaS – System Infrastructure Software market segments) in 2021.

[21] 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.

A6.9 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 A6.2 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 A6.2: 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.[22]
AWS	✓	✓	✗[23]	AWS provided a breakdown of AWS costs in more detail than that available publicly.[24]
Microsoft	✓	Some	✗	Microsoft provided some profit information for Azure, but this did not include an allocation of all operating costs.[25]
IBM	✓	Some	✗	IBM commented on the operating profit for its public cloud business and provided only high-level estimates.[26]
Oracle	✓	✗	✗	Oracle said it could not provide profit information on cloud beyond its public reporting for the Cloud and License segment.[27]

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.

[22] Google response dated 16 December 2022 to our 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.
[23] We note that Amazon does publish some asset information for AWS in its financial statements, as discussed further below.
[24] 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).
[25] 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).
[26] 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.
[27] Oracle response dated 16 December 2022 to our 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.

Hyperscaler businesses we focus on for profitability analysis

A6.10 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:

- **Amazon**: Our analysis references **AWS** as it provides cloud infrastructure services. The information available is largely sufficient for our analysis.
- **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 2022 it represented around [✂] of Microsoft Cloud revenue[28] - 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.
- **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 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 [✂].[29] We consider that Workspace is likely to be a smaller part of Google Cloud than Office 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

A6.11 Of the companies featured in Table A6.1, 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.[30]

[28] 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).

[29] Google response dated 16 December 2022 to our s.174 notice dated 26 October 2022, Part B question 4.

[30] 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 used annual average exchange rates from S&P Capital IQ Pro to convert into pound sterling.

Hyperscaler cloud profitability

A6.12 In this section we set out the hyperscalers cloud revenues and profits and compare these to other cloud providers. We consider the following:

- Relative contribution of cloud to hyperscalers' overall businesses
- Trends in hyperscaler cloud revenues
- Trends in hyperscaler cloud profit margins

The relative contribution of cloud to the hyperscalers' overall businesses

A6.13 Table A6.3 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).

Table A6.3: 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		
- Azure	*[✂]*	*n/a*
- Microsoft Cloud	*46%*	*n/a*

Source: Ofcom analysis of cloud providers' published financial statements for financial years ending December 2022 (Amazon and Alphabet) and June 2022 (Microsoft) and information provided by Microsoft in response to our information requests.[31]

A6.14 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[32] across Amazon's other reported segments.[33] In contrast, Microsoft Cloud represented almost half of Microsoft's revenue in its 2022 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 [✂].[34] Google Cloud is a relatively small yet growing contributor to Alphabet's revenue and is currently making operating losses.

[31] 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).
[32] Operating profit (EBIT) margin is calculated as operating profit divided by revenue.
[33] 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.
[34] 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).

Trends in hyperscaler cloud revenues

A6.15 Figure A6.4 below shows publicly reported global revenues associated with cloud activities for the three hyperscalers plus five global competitors (IBM, 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 section.

Figure A6.4: Annual global revenue for major cloud providers from their latest reported financial year (£bn)

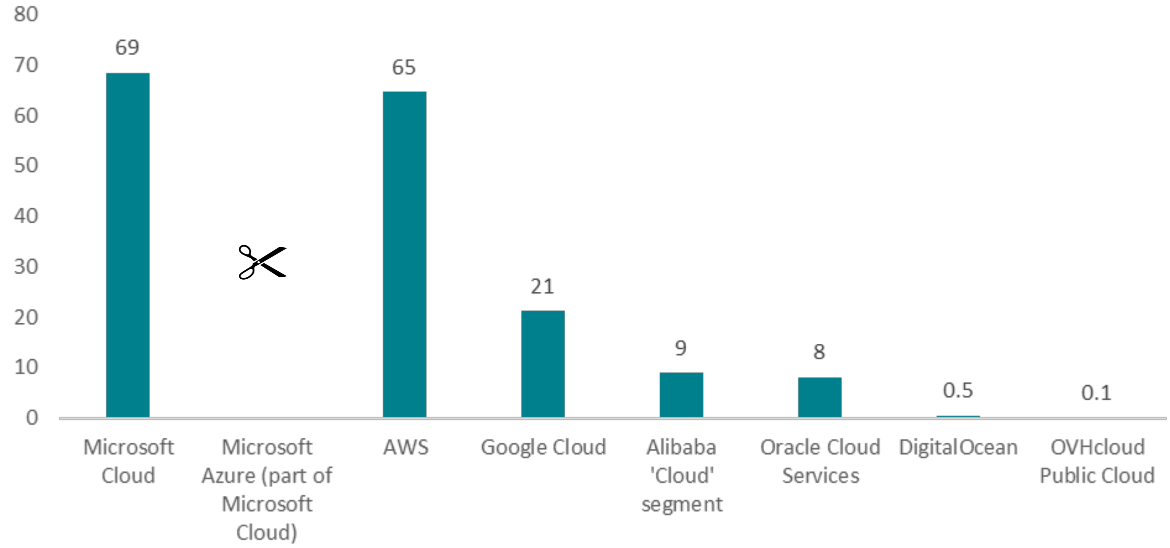

Source: Ofcom analysis of cloud providers' published financial statements and information provided by Microsoft in response to our information requests.[35] Figures (other than Azure) come from the following financial statements: Microsoft – year to June 2022, AWS – year to December 2022, Google Cloud – year to December 2022, Oracle Cloud Services – year to May 2022 (data on 'cloud services' taken from quarterly earnings releases), Alibaba 'Cloud' segment – year to March 2022, DigitalOcean – year to December 2022, OVHcloud Public Cloud – year to August 2022.

A6.16 IBM reports 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.[36] 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 2021 its global public cloud revenue was [✂], which was lower than the hyperscalers.[37]

---

[35] 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).

[36] 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*".

[37] IBM response dated 6 December 2022 to our s.174 notice dated 25 October 2022, Part B question 4 (Annex 2 submission).

A6.17 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, Azure represents around [✂] of Microsoft Cloud revenues and AWS revenues. Google Cloud's most recent reported revenues were closer to those of smaller cloud providers like IBM and Oracle than AWS.

A6.18 Figure A6.5 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 2022.

Figure A6.5: Quarterly year-on-year global revenue growth for major cloud providers, 2018-22

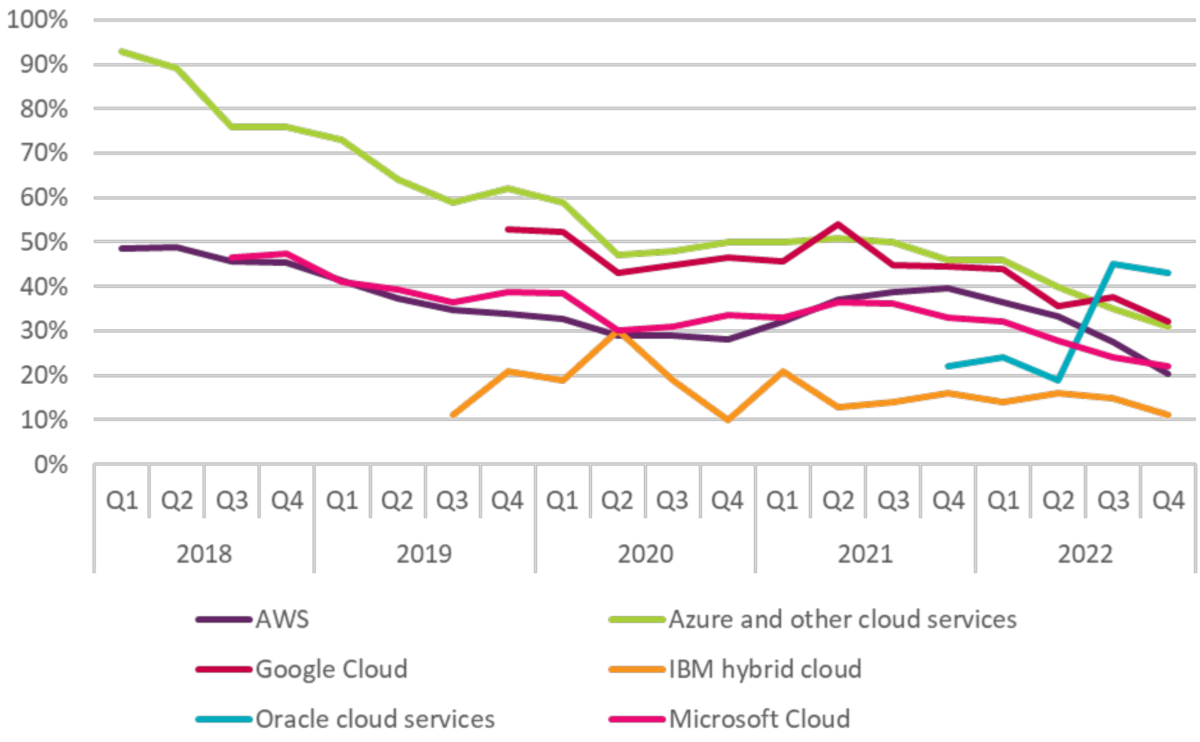

Source: Ofcom analysis of cloud providers' published financial statements.

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. 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.

A6.19 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.

A6.20 The chart also indicates that, over this period, Azure's revenue grew faster than Microsoft Cloud's revenue, and Microsoft's cloud activities grew quicker than other cloud providers.

A6.21 We also obtained UK revenue estimates from cloud providers through our statutory information requests, which cover the period 2019 to 2021. This data suggests that during this period UK revenue growth was broadly consistent with global revenue growth.

Drivers of cloud revenue growth

A6.22 Growth in cloud revenues could be caused by:

- Increases in the unit prices customers are charged for consuming cloud services
- Growth in the number of customers purchasing cloud services
- Growth in the usage of chargeable services consumed by existing customers

A6.23 Based on responses to our statutory information requests, revenue growth appears to be, at least in part, due to more customers[38], but we cannot determine whether some revenue growth was because of price increases or more usage by existing customers.

Trends in cloud profit margins

A6.24 In this section we compare the operating profit (EBIT[39]) 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.

Estimating Microsoft Cloud EBIT

A6.25 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.[40] Microsoft Cloud gross margins have been around 70% in Microsoft's last three financial years. As Microsoft also reports revenue for Microsoft Cloud, we estimated the gross margin (in absolute terms) by multiplying revenue by the gross margin percentage.[41]

---

[38] 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 our s.174 notice dated 24 October 2022, Part B question 15; Google response dated 16 December 2022 to our s.174 notice dated 26 October 2022, Part B question 15; Microsoft response dated 9 December 2022 to our s.174 notice dated 21 October 2022, Part B question 15.

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

[40] 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 65 of its 2022 10-K.

[41] 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

A6.26 To estimate EBIT for the financial years ending 2016 to 2022, 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 90% of Microsoft's operating costs in its 2022 financial year, with General and Administrative costs representing around 10%.[42]

A6.27 There is uncertainty about what proportion of these costs relate to Microsoft Cloud. Microsoft's 10-K[43] 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.[44]

A6.28 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.[45] As these represent the majority of operating costs, we used revenue as the basis of allocation.[46]

Estimating Azure EBIT

A6.29 Microsoft provided us with an estimate of EBIT for Azure for its financial years ending 2018 to 2022 in response to our statutory information requests.[47] These estimates included an allocation of infrastructure costs to Azure but did not include an allocation of other operating costs.

A6.30 Microsoft also provided us with a full 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.[48]

A6.31 We asked Microsoft to explain how costs are allocated to its cloud businesses, but it did not explain its allocation approach in detail.[49] We therefore needed to make some

---

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.

[42] Calculated based on figures presented on page 57 of Microsoft's 2022 10-K.

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

[44] Page 94 of Microsoft's 2022 10-K.

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

[46] 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 16 of Microsoft's 2022 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.

[47] 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).

[48] 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.

[49] 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.

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 total Microsoft operating costs for Cloud and Enterprise is very similar to its share of total Microsoft revenue.[50] This suggests that allocating operating costs to Azure using revenue could give a reasonable estimate of its EBIT, given that Azure is part of Cloud and Enterprise.

A6.32 As for Microsoft Cloud, we also considered allocating 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.

Comparison of EBIT with other cloud providers

A6.33 Figure A6.6 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.[51]

A6.34 This chart shows that in absolute terms, the most recent annual EBIT for AWS and our estimate of EBIT for Microsoft Cloud are significantly higher than the EBIT of other cloud providers, including Google Cloud, which is loss making. We estimate that Azure EBIT in Microsoft's 2022 financial year was [✂]. This is [✂].

[50] 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.
[51] 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.

Figure A6.6: Annual global EBIT for the latest financial year (£bn)

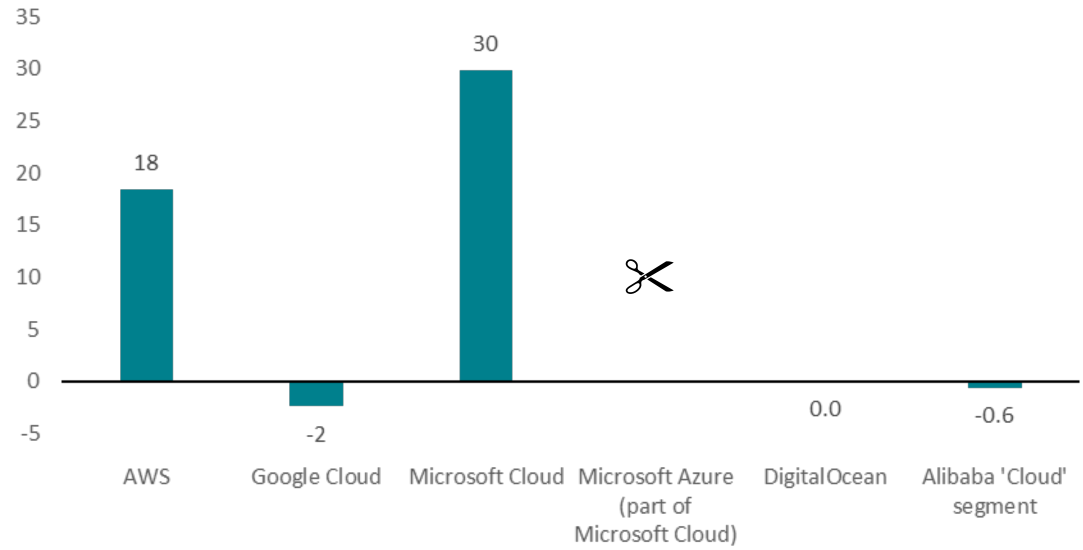

Source: Ofcom analysis of cloud providers' published financial statements, information provided by Microsoft in response to our information requests[52] and Ofcom assumptions.

A6.35 Figure A6.7 below captures quarterly EBIT margins between 2018 and 2022 for AWS, Google Cloud, Alibaba's 'Cloud' segment and DigitalOcean.[53] 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[54] and an annual EBIT margin for Azure). Our estimates for Azure have been redacted in the published version of this report.

[52] 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.

[53] Operating profit (EBIT) margin is calculated as operating profit divided by revenue.

[54] We are able to 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 A6.7: Quarterly EBIT margins for major cloud providers, 2018 to 2022

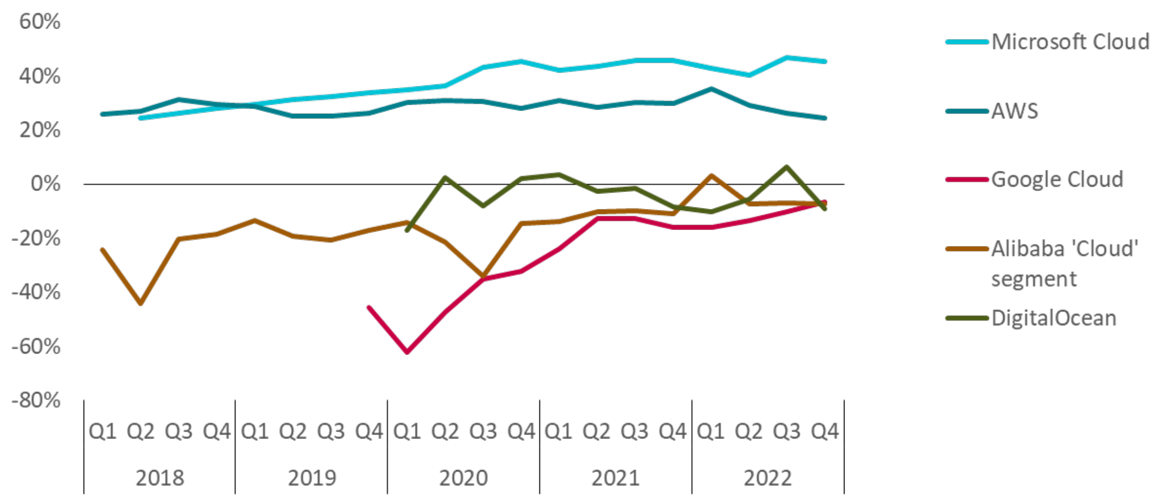

Source: Ofcom analysis of cloud providers' financial data reported by S&P Capital IQ, information provided by Microsoft in response to our information requests[55] and Ofcom assumptions. We have used the latest filings (incorporating restatements).

A6.36 OVHcloud's 'Public Cloud' segment had an EBITDA margin[56] of 35% to 40% in its 2021 and 2022 financial years respectively.[57] 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.[58]

A6.37 For the hyperscalers, we estimate that, over this period:

- AWS had stable EBIT margins of around 20% to 30%.
- 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 [✂].
- Google Cloud was loss making despite strong revenue growth, though quarterly losses steadily reduced. It appears that Google Cloud's current losses are driven by the need to invest ahead of revenues to catch up with other hyperscalers, particularly in relation to its salesforce.[59] We also observe that when AWS was generating similar revenues to Google Cloud, its reported EBIT margins were around 30%.[60] This suggests Google

---

[55] 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.
[56] EBITDA stands for Earnings Before Interest, Tax, Depreciation and Amortisation.
[57] Calculated from OVH Cloud 2022 Financial Statements, page 27.
[58] 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.
[59] For example, page 11 of Alphabet's 2022 10-K, says Google is "*incurring costs to build and maintain infrastructure to support cloud computing services, invest in cybersecurity, and hire talent, particularly to support and scale our sales force".* Also, in its Q4 2022 earnings call, Alphabet's CFO said Google Cloud is investing ahead of revenues to help support its customers, but that it is focused on the path to profitability.
[60] In Q4 2022 Google Cloud revenue was $7.32bn, as reported in page 2 of Alphabet's Q4 2022 results press release. Our analysis of quarterly AWS revenue shows that AWS revenue was closest to this in Q4 2018 when AWS generated revenue of $7.43bn, as reported in page 10 of Amazon's Q4 2018 quarterly earnings release. Combined with the reported AWS operating income of $2.18bn, this tells us that AWS EBIT margin was 29.3% in Q4 2018.

Cloud's current losses may not be due to lower revenue alone, as its costs are higher than AWS's when it was in a similar position.

A6.38 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).

A6.39 We also received additional insights from responses to our statutory information requests.

- 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.[61]
- 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.[62]

A6.40 Overall, this evidence indicates that cloud profits for AWS and Microsoft Cloud are higher than other cloud providers, for whom, in some cases, operating losses to date have been common. While there is some evidence of improving profits (or reduced losses) among smaller cloud providers, these are a lot lower than those for AWS and Microsoft Cloud. Azure operating profits [✂].

Assessment of ROCE relative to WACC

Why we use ROCE as a measure of profitability

A6.41 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[63] and has been frequently used in past CMA market studies and Ofcom profitability assessments.[64]

A6.42 ROCE can be compared against the weighted average cost of capital (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 costs will have high margins, but that would not

[61] 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.

[62] 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.

[63] Competition Commission (2013). Guidelines for market investigations: Their role, procedures, assessment and remedies, Annex A paragraph 9 [accessed 28 March 2023].

[64] See for example: (1) CMA, 2022. Appendix D: financial analysis of Apple's and Google's mobile ecosystems [accessed 28 March 2023]; (2) Ofcom, 2022. Annex: Ofcom's future approach to mobile markets [accessed 28 March 2023].

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.[65]

A6.43 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.

A6.44 The CMA's guidance therefore says a finding that 'profitability of firms which represent 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'.[66] 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.

Our approach to estimating ROCE

A6.45 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.[67]

A6.46 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.

A6.47 Capital employed should represent a reasonable estimate of what it would cost for a competitor to replicate the operational assets and liabilities of the relevant business. 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:

- **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

[65] CMA, 2020. Appendix D: profitability of Google and Facebook, paragraphs 8 and 9 [accessed 28 March 2023].
[66] Competition Commission, 2013. Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraph 118 [accessed 28 March 2023].
[67] 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 28 March 2023].

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, 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.[68]

- **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 relevant business.[69] 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

A6.48 Our ROCE analysis focuses on AWS and Microsoft Azure, as our shares of supply analysis indicates they represent a substantial share of cloud infrastructure revenues, and we know that Google Cloud is currently loss-making and will have a negative ROCE. For context and comparison with Azure, we also present an estimate of ROCE for Microsoft Cloud.

A6.49 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 and Azure 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. We have therefore had to make assumptions when estimating capital employed, as set out below.

A6.50 For each of Amazon and Microsoft, we set out below how we have estimated ROCE.

A6.51 The rest of this section is structured as follows:

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

[68] 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 28 March 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 28 March 2023].

[69] 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.

WACC applicable to cloud activities

A6.52 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.

A6.53 We asked AWS and Microsoft if they estimated WACC for AWS and Azure, but both said they did not produce WACC estimates at this level.[70] We have therefore estimated a WACC that we think captures the systematic risk associated with cloud activities.

A6.54 We think the pre-tax nominal WACC applicable to AWS and Azure for the period we are considering is likely to fall between 9.0% and 13.0% for the following reasons:

- 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.[71]
- When estimating the WACC for Google and Facebook, the CMA used an equity beta range of 1.0 to 1.15. This range appears to capture recent equity beta estimates for the hyperscalers[72] and cloud computing exchange traded funds.[73] However, we observe that equity betas for the hyperscalers were somewhat higher before Covid lockdowns in 2020. For example, Amazon and Microsoft equity betas were occasionally between 1.2 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 equity beta range of 1.0 to 1.30 to reflect the possibility that cloud betas could have been higher in the past.[74]
- 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%.
- Combining our equity 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%.[75] 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.

---

[70] AWS response dated 9 December 2022 to our s.174 notice dated 24 October 2022, Part B question 29; Microsoft response dated 9 December 2022 to our s.174 notice dated 21 October 2022, Part B question 29. We note we did not ask Microsoft if it estimated WACC for Microsoft Cloud.
[71] CMA, 2020. Appendix D: profitability of Google and Facebook, pages D39 to D41 [accessed 28 March 2023].
[72] The hyperscalers' betas will reflect the systematic risk of their overall businesses, not just cloud, though cloud represents an increasing proportion of overall activity.
[73] 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.
[74] 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.
[75] 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.

A6.55 As a sense check, we 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[76], 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%[77] and in subsequent assessments it generally reduced, falling to around 10% in our most recent decision in 2021.[78] On this basis, our estimated range of a 9% to 13% for cloud does not appear unreasonable.

Our assessment of ROCE for AWS

A6.56 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.

A6.57 To derive our baseline ROCE estimates for AWS, we used AWS's reported EBIT.

A6.58 To derive capital employed for AWS we considered the information Amazon publishes in its annual reports:

- **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.[79] This category also includes finance leases (which could include data centres owned under leasing arrangements).[80]
- **AWS total assets**: Amazon says AWS total assets primarily consist of property and equipment assets and accounts receivable.[81] Amazon also says that the majority of technology infrastructure assets are allocated to AWS based on usage.[82] 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.[83]

A6.59 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 a competitor may need to replicate. While this measure does not include operating leases, Amazon's public reporting implies that these may be more relevant to its retail operations

[76] BT's ICT operations are spread across its BT Global and other divisions. These include global security, cloud and networking services.
[77] Ofcom, 2016. Business Connectivity Market Review 2016, Annex 30, Table A30.2 [accessed 28 March 2023].
[78] Ofcom, 2021. Wholesale Fixed Telecoms Market Review 2021-26, Annex 21, Table A21.2 [accessed 28 March 2023].
[79] Amazon, 2023. 2022 10-K, page 47 [accessed 28 March 2023].
[80] Amazon, 2023. 2022 10-K, page 47 [accessed 28 March 2023].
[81] Amazon, 2023. 2022 10-K, page 68 [accessed 28 March 2023].
[82] Amazon, 2023. 2022 10-K, page 68 [accessed 28 March 2023]. We understand technology infrastructure includes servers, networking equipment and data centres.
[83] Amazon, 2023. 2022 10-K, page 68 [accessed 28 March 2023].

than AWS.[84] 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)[85], 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.

A6.60 Figure A6.8 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.

Figure A6.8: AWS global EBIT and net property and equipment assets, 2013 – 2022, $bn

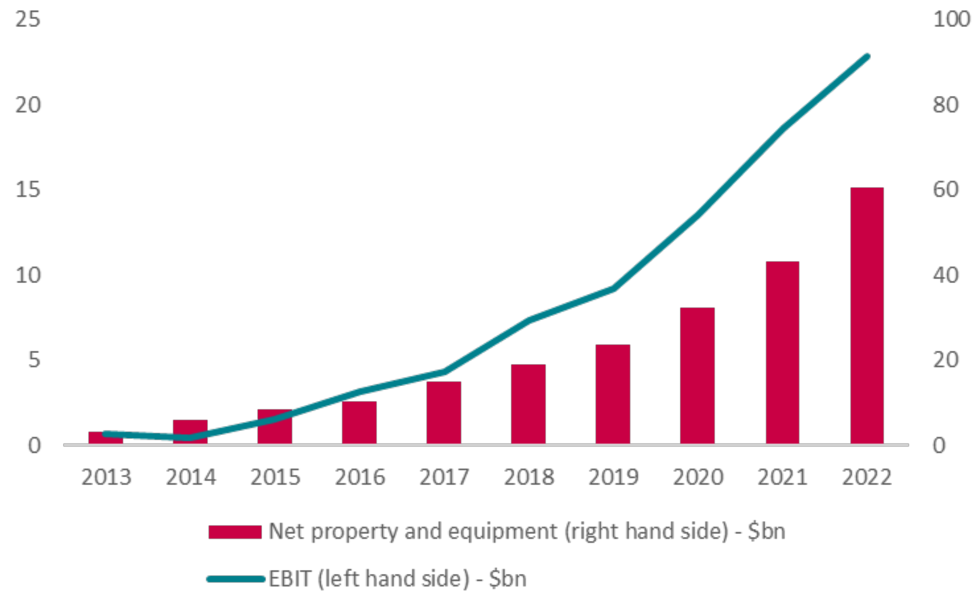

Source: Ofcom analysis based on Amazon 10-K reports.

A6.61 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:

- **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

[84] 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.
[85] 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).

accounts receivable but not accounts payable (or any other relevant current liabilities). It may also include some assets which a competitor would not require to operate AWS. While we do not have the data to adjust the asset base to remove assets that a competitor may not require, 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.

- **Sensitivity 2: Include all technology infrastructure assets with AWS capital employed.** Amazon allocates technology infrastructure assets to AWS on the basis of 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.[86] As a single data centre could be used by both external and internal customers, we have considered the effect of assuming that a competitor may need to replicate all of these technology assets 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, a competitor may replace these assets with smaller data centres or fewer servers to reflect the fact only external customers were being served. We are only able to make this adjustment for the period 2017 to 2021, based on information provided by AWS.[87]

A6.62 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. We do not have visibility of the value of these assets for Amazon overall or for AWS, so we have not made an adjustment for this. However, based on the results of our wider analysis we consider it is unlikely that any such adjustments would be substantial enough to affect our overall analysis.

A6.63 Figure A6.9 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).

A6.64 Our baseline estimate indicates that AWS ROCE has increased since 2013 and has been around 40% since 2018. It is higher than our estimate of WACC in all years except 2014, i.e. in nine of the last ten years.

A6.65 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 baseline ROCE, the ROCE under this sensitivity is above our estimate of WACC in all years except 2014.[88]

[86] 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.
[87] 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. We were also unable to adjust depreciation associated with technology infrastructure assets, though we do not think doing so would significantly affect the results for this sensitivity.
[88] 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.

Figure A6.9: AWS ROCE estimates (baseline approach and first sensitivity), 2013-22

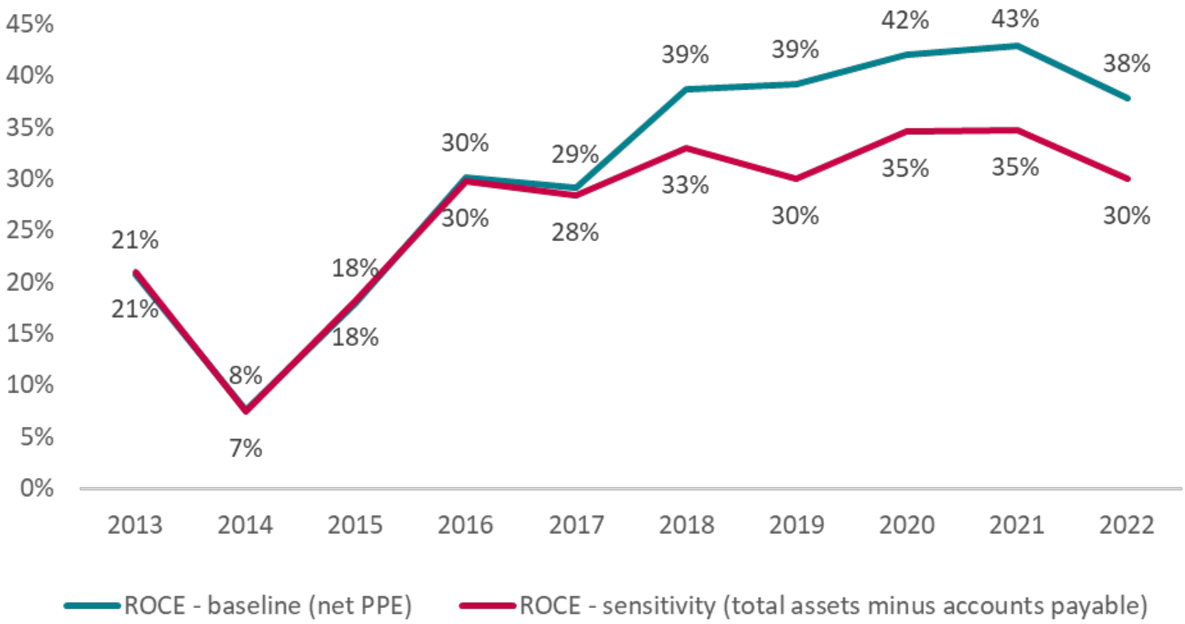

Source: Ofcom analysis based on public information from Amazon 10-K reports and Ofcom assumptions.

A6.66 Figure A6.10 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 2021.

A6.67 Our ROCE estimates under this sensitivity [✂].

Figure A6.10: AWS ROCE estimates (addition of non-AWS technology infrastructure assets), 2017-21

[✂]

Source: Ofcom analysis based on public information from Amazon 10-K reports and data provided by AWS.[89]

Our assessment of ROCE for Microsoft Azure

A6.68 To derive our baseline ROCE estimates for Azure, we used our estimated EBIT for Azure. We explained how we estimated this in the previous section.

A6.69 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.[90] We have therefore estimated capital employed for Azure as follows.

[89] 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.

[90] Microsoft response dated 9 December 2022 to our 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.

A6.70 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 2022. It also provided the same information for Microsoft overall.

A6.71 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.

A6.72 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[91] and ii) 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, [✂][92], indicating assets associated with these costs are likely to form a significant proportion of Azure capital employed.

A6.73 The ratio we applied to Azure's infrastructure operating costs ranges from [✂] for Microsoft's financial years ending 2018 to 2022.[93] 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.[94]

A6.74 Our ROCE estimates for Azure are shown in the figure below.

Figure A6.11: Azure ROCE estimate, Microsoft financial years ending 2018 - 2022

[✂]

Source: Ofcom analysis based on Microsoft 10-K reports and information provided by Microsoft in response to our information requests.[95]

A6.75 We estimate that Azure's ROCE increased over this period and is now above our estimate of WACC. [✂].

A6.76 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

[91] After the exclusion of goodwill from non-current assets, consistent with our approach to capital employed set out above. In the financial year ended June 2022, Microsoft reported net property and equipment of $74.4bn, operating lease assets of $13.1bn, total non-current assets of $195.2bn and goodwill of $67.5bn. These figures are reported in page 59 of Microsoft's 2022 10-K. This means that net property and equipment plus operating leases together accounted for 68.6% of total non-current assets (excluding goodwill).
[92] 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).
[93] 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).
[94] For example, if the mix of assets associated with Azure had a longer useful life than for Microsoft overall.
[95] 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.

provide us with all the financial information we requested on Azure. However, our analysis is [✂].[96]

Our assessment of ROCE for Microsoft Cloud

A6.77 To derive our baseline ROCE estimates for Microsoft Cloud, we used our estimated EBIT for Microsoft Cloud. We explained how we estimated this above.

A6.78 Microsoft does not separately report assets for Microsoft Cloud.[97] As a result we estimated capital employed for Microsoft Cloud by reference to capital employed for Microsoft's overall business. 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 2022.

A6.79 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 a competitor will likely need to replicate these types of assets to provide a service similar to Microsoft Cloud. This approach broadly aligns with our baseline approach to calculating AWS ROCE above.

A6.80 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).[98]

A6.81 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 2022 (Microsoft's financial year ends in June).

[96] 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 our 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.
[97] We did not ask Microsoft to provide us with balance sheet information on Microsoft Cloud specifically, but its 10-K reporting says that assets are not allocated to segments for internal reporting presentations (see page 95 of its 2022 10-K).
[98] 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 and other 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 A6.12: Baseline ROCE and sensitivity ROCE estimates for Microsoft Cloud, Microsoft financial years ending 2016-22

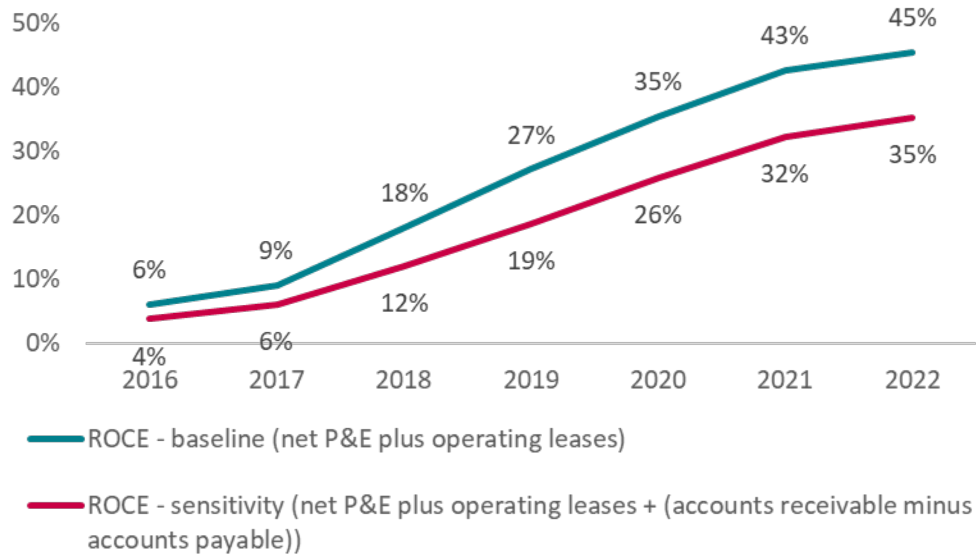

Source: Ofcom analysis based on public information from Microsoft 10-K reports and Ofcom assumptions.

A6.82 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 45% in Microsoft's financial years ending 2018 to 2022, above our estimate of WACC.

A6.83 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.

A7. Multi-cloud and switching

A7.1 In this annex we set out further evidence on the prevalence of multi-cloud and switching. As part of this, we present findings from the Context Consulting market research and compare them with other industry reports, drawing out differences and possible explanations, where we identify them.

A7.2 As set out in Section 3, there are several ways to categorise the different forms that multi-cloud can take. The categorisation we have adopted for this interim report has been developed based on the evidence we have gathered and assessed so far, and we consider it useful for our purpose of assessing the state of competition in cloud infrastructure services. Meanwhile, the Context Consulting market research was commissioned at an earlier stage of the market study process when our categorisation of the different forms of multi-cloud was not fully aligned with the one presented in Section 3. As a result, in some cases 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.

A7.3 We also note that Context Consulting market research asked respondents about switching and multi-cloud in relation to IaaS/PaaS providers. Customers may have understood such providers to include both cloud providers and ISVs. This means responses regarding the use of multi-cloud are likely to include cases where a customer is using more than one supplier of cloud infrastructure services (including cloud providers and ISVs) on the same cloud 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.

A7.4 To capture a richer and more in-depth understanding of customer responses on switching and multi-cloud we commissioned a further piece of qualitative market research to Context Consulting. Overall, we consider that the market research (including the quantitative and qualitative stages) provides useful evidence, which we set out below, on multi cloud albeit with some caveats.

Evidence on multi-cloud in practice

Prevalence of multi-cloud

A7.5 About 52% of IaaS and/or PaaS users in the market research use more than 1 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%.

A7.6 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 interim report). This is because 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 cloud providers, the 52% figure may capture: (i) customers using cloud providers and ISVs hosted on the same cloud; (ii) customers using a combination of private and public cloud providers (i.e. hybrid cloud), or (iii) customers using multiple private cloud providers. The 61% figure would exclude group (iii) but may still include customers from groups (i) and (ii).

A7.7 The market research also asked respondents who use only 1 IaaS/PaaS provider and those considering purchasing IaaS/PaaS[99] about their attitude to **potentially using multi-cloud architecture in the future**. A majority responded, "yes definitely" (23%) or "yes possibly" (63%).[100]

A7.8 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).[101] This may be explained by the fact that the Context Consulting market research was aimed at IaaS and/or PaaS use only, and we did not ask about SaaS providers. 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.

A7.9 A 2023 report commissioned by Oracle suggests that 98% of respondents were using (or planning to use within 6 months) more than 1 IaaS/PaaS provider.[102]

A7.10 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).[103]

Management of multi-cloud architecture

A7.11 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.[104]

A7.12 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

[99] 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.
[100] Context Consulting market research report, slide 79.
[101] Flexera, 2023. 2023 State of the Cloud report, page 18 [accessed 30 March 2023].
[102] Multicloud in the Mainstream: Making IT Work 'As Advertised'. Commissioned by Oracle. S&P Global Market Intelligence. February 2023. Figure 1, p.1.
[103] https://resources.foundryco.com/download/cloud-computing-executive-summary, page 3 [accessed on 13.02.2023].
[104] Context Consulting market research report, slide 79.

(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.[105]

A7.13 In addition, as discussed 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 cloud providers (i.e. the definition of multi-cloud used in this interim report), these results may capture cases of customers using a cloud provider and an ISV hosted on the same cloud, 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.

A7.14 The 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,[106] 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.

A7.15 In the course of this market study we issued a customer questionnaire to several companies, asking them about their multi-cloud use.

A7.16 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.[107] Where the [✂] adopted

[105] 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.
[106] We consider this figure is likely to overstate prevalence of multi-cloud. 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 cloud providers (i.e. the definition of multi-cloud used in this interim report), this result may capture cases of customers using a cloud provider and an ISV hosted on the same cloud, 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.
[107] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].

this approach, it is generally for workloads which use a single service at very large scale (where the additional time and costs are not material), and for the [✂] these are [✂].[108]

A7.17 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.[109] 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.

A7.18 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].[110] 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)].[111]

A7.19 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.[112] 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.[113] This customer suggested there are significant advantages from having all their cloud services with a single cloud provider:

- The ability to devote resources to recruiting, training and retraining a dedicated team, specialising in architecture of just one cloud provider.
- The integration cost and effort is reduced and there are administrative advantages from having a single billing system.
- The ability to qualify for discounts or, more specifically, to meet usage commitments.[114]

A7.20 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.[115]

A7.21 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

[108] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[109] [✂].
[110] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].
[111] [✂] response dated [✂] to the s.174 notice dated [✂] question [✂].
[112] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].
[113] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[114] [✂] response dated [✂] to the s.174 notice dated [✂], questions [✂].
[115] Context Consulting market research report, slide 80.

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”[116] is now at 42%.[117]

Figure A7.1: Use of multi-cloud in the Flexera 2023 State of the Cloud report[118]

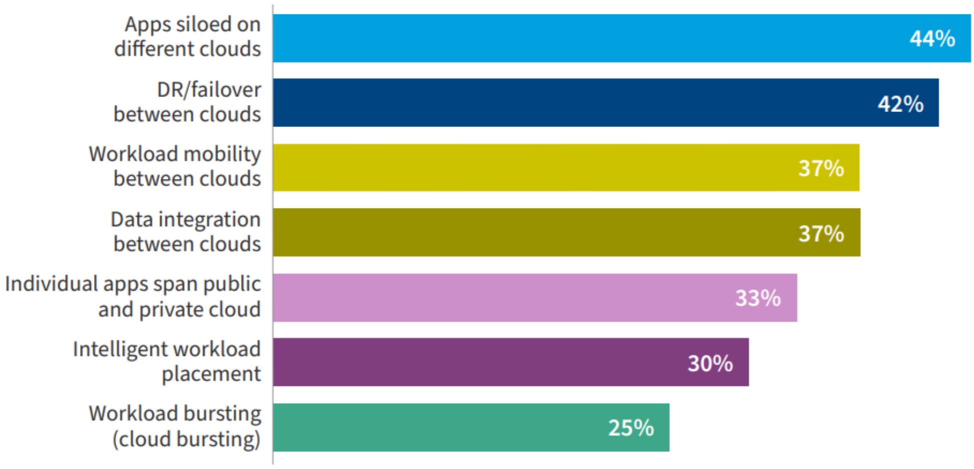

Source: Flexera 2023, State of the Cloud report, page 20.

Benefits and challenges of using multi-cloud

A7.22 All respondents to the market research (those already using more than 1 IaaS/PaaS provider, those using just 1, 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.[119] Figure A7.2 below shows that respondents who already use more than 1 IaaS/PaaS provider tend to report benefits more often than those who only use 1 provider.

[116] DR is disaster recovery.
[117] Flexera, 2023. 2023 State of the Cloud report, page 20 [accessed 30 March 2023].
[118] Proportion of all respondents.
[119] Context Consulting market research data tables, Q32.

Figure A7.2. Benefits of using multi-cloud, by the number of providers respondents use

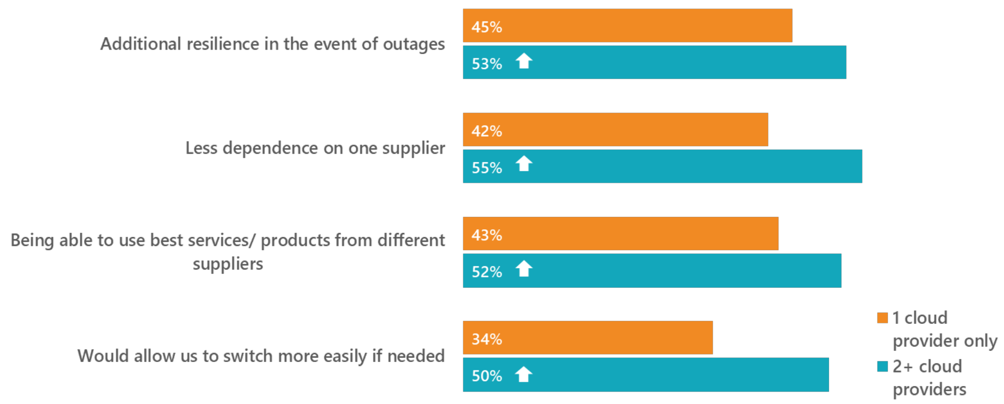

Source: Context Consulting market research report, slide 78.

A7.23 In the market research, 46% of those who used 2 IaaS/PaaS providers, and 52% of those who used 3 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.[120]

A7.24 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%).[121] These results are broadly consistent with what we saw in the market research.

A7.25 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.[122] The report also categorises the drivers for multi-cloud adoption into two broad groupings:

- Defensive — i.e. intended to guard against pricing or total-cost-of-ownership shifts, vendor lock-in, regulatory compliance and business continuity interruptions.
- 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.

[120] Context Consulting market research data tables, Q64.
[121] https://resources.foundryco.com/download/cloud-computing-executive-summary, page 3 [accessed on 13.02.2023].
[122] 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 A7.3. Benefits of using multi-cloud highlighted by the 2023 Oracle report

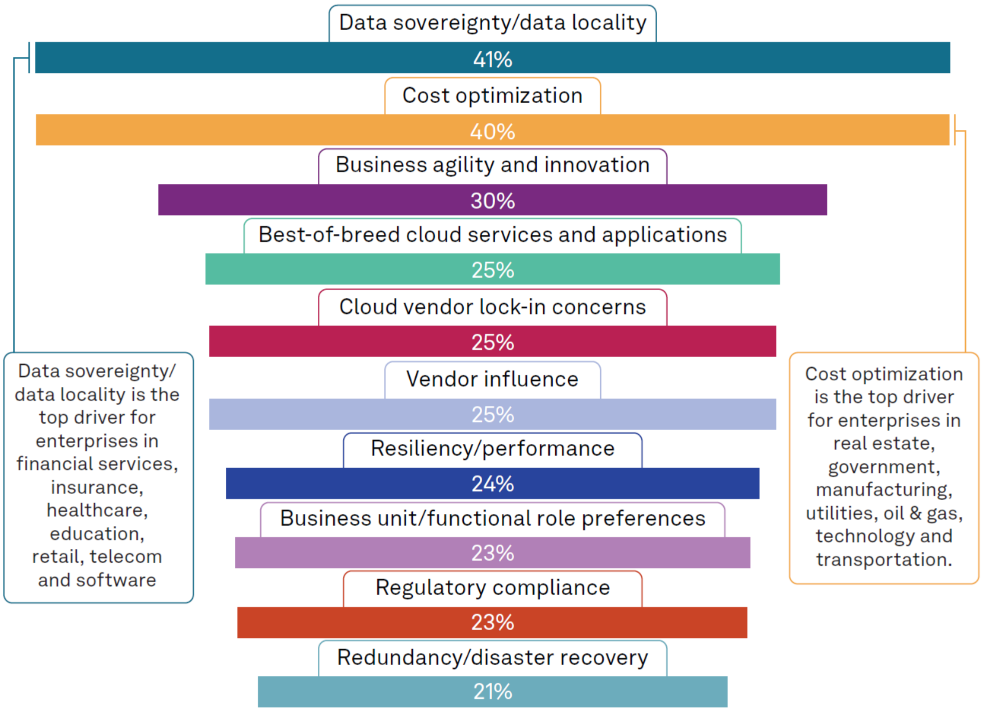

Source: Multicloud in the Mainstream: Making IT Work 'As Advertised'. Commissioned by Oracle. S&P Global Market Intelligence. February 2023.

A7.26 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 A7.4: Challenges of using multi-cloud as reported in the Context Consulting market research

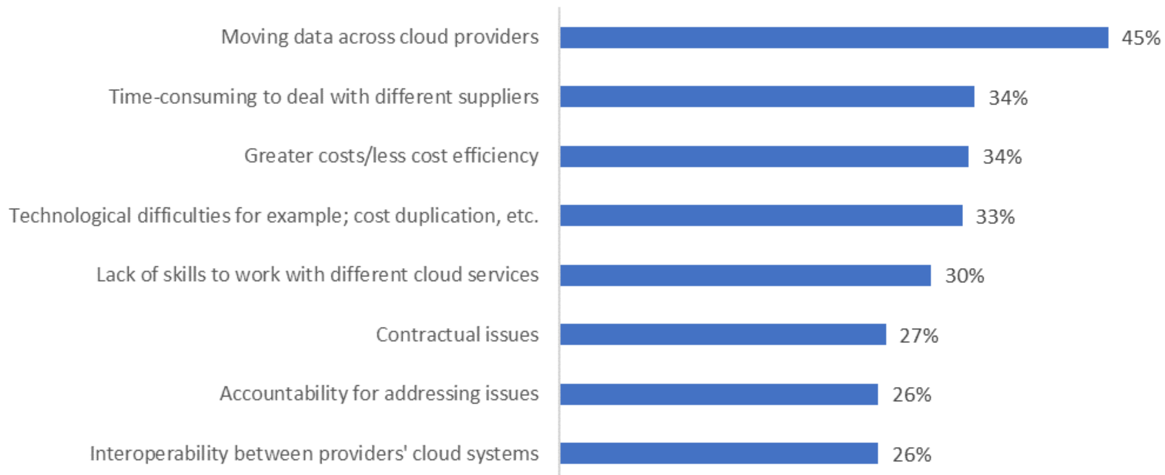

Source: Ofcom analysis of Context Consulting market research data tables, question 31.

A7.27 Comparing responses of those who already use more than one IaaS/PaaS provider with those who only use one likely 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 A7.5 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 A7.5: Challenges of using multi-cloud as reported in the Context Consulting market research

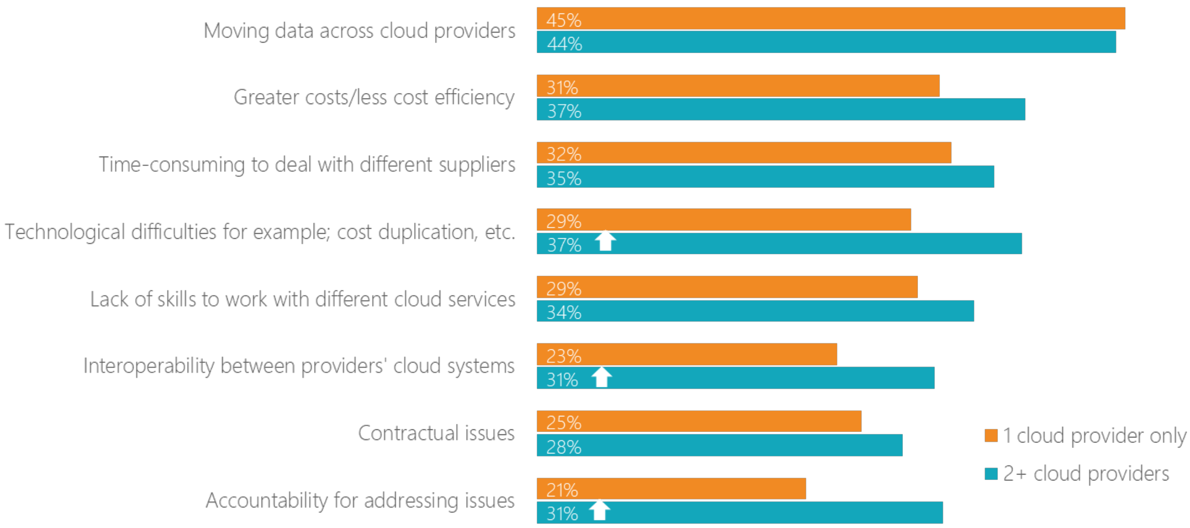

Source: Context Consulting market research report, slide 79.

A7.28 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).

A7.29 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 this issue essentially comes down to shortage of talent and expertise. [123]

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

Figure A7.6: Challenges of multi-cloud according to the 2023 Oracle report

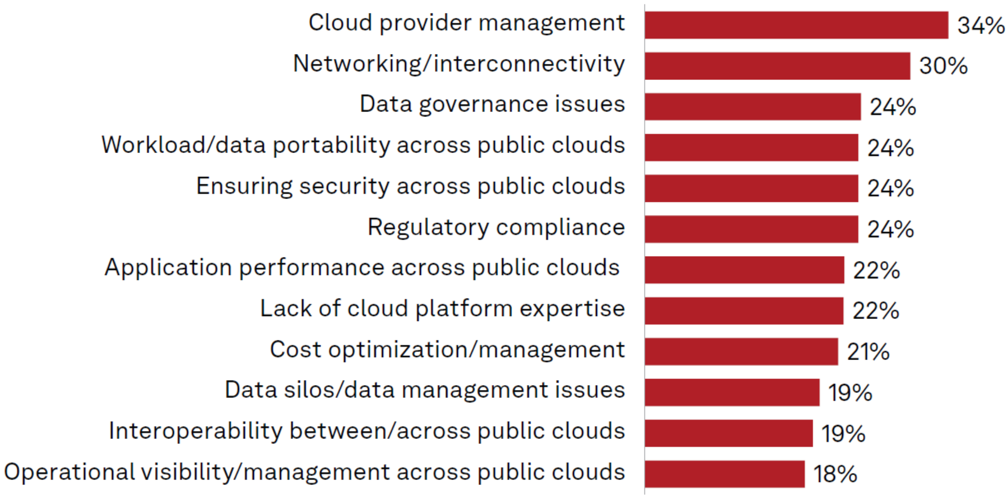

Source: Multicloud in the Mainstream: Making IT Work 'As Advertised'. Commissioned by Oracle. S&P Global Market Intelligence. February 2023.

A7.30 We also asked some customers to identify obstacles to 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.[124] 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.[125]

A7.31 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.

Evidence on switching in practice

Prevalence of switching

A7.32 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.

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

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.

A7.33 There are differences in the proportions of respondents who switched according to their attitude to technology adoption (see Figure A7.7). 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 A7.7: Switching behaviour in the Context Consulting market research by attitude to technology

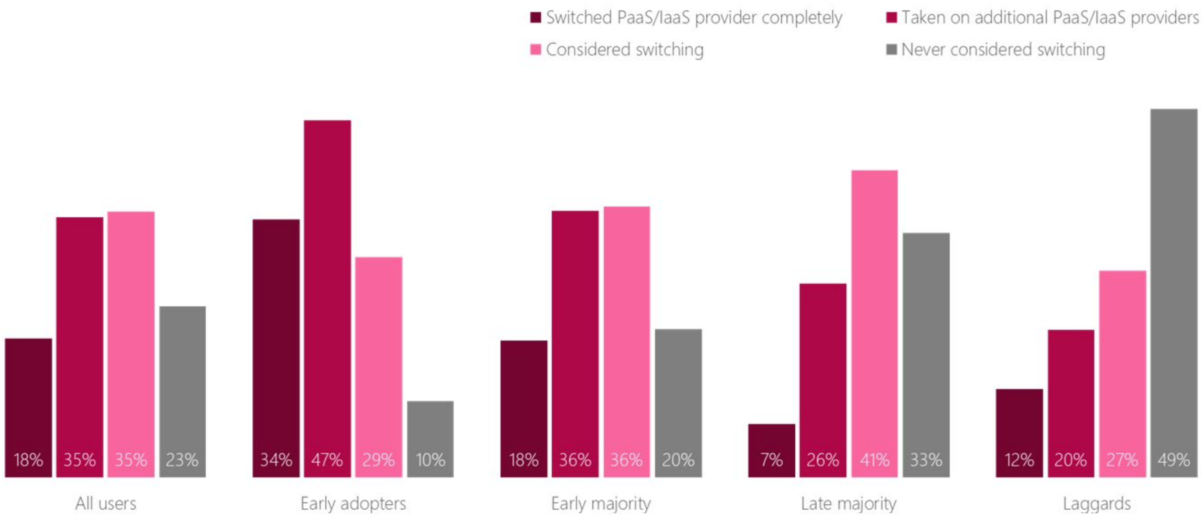

Source: Context Consulting research report, slide 105.

A7.34 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.

A7.35 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%). Some of these categories probably overlap with the "laggards"[126] 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.

A7.36 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.[127] This number is similar to our finding of 18% of respondents who switched a IaaS/PaaS provider.

A7.37 The qualitative part of the Context Consulting market research, 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:

- We encountered few if any examples of organisations switching away from one of the hyperscalers.
- The switching we found was typically from a data centre to an IaaS/PaaS environment.
- In some cases, firms were adding additional platforms (e.g. bringing AWS onboard alongside Azure).
- It was still relatively early in the adoption journey for most companies, and they were evaluating progress rather than looking to make significant changes.
- In most cases, firms were still on the way in, not out, of their IaaS/PaaS environments.[128]

A7.38 Overall, it is possible that the respondents to the quantitative part of the market research were conflating a migration from on-premise IT/private cloud to public cloud, switching between first- and third-party services within a public cloud and switching between public clouds.

A7.39 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 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.

[126] 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.
[127] Report on Trade Practices in Cloud Services Sector. June 2022. Available at https://www.jftc.go.jp/en/pressreleases/yearly-2022/June/221102EN.pdf (accessed on 15.02.2023). Page 49.
[128] Context Consulting market research report, slide 110.

A7.40 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.[129] Several respondents moved significant workloads from one cloud provider to another (i.e. partial switch) which indicates that customers are avoiding a full switch and are only looking to switch where necessary. [✂][130]

A7.41 Some customers ([✂]) 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.[131] Some commented that when a new workload appears, they consider to which of their several cloud providers they could best allocate it, suggesting this does not require moving all workloads.

Reasons for switching and switching experience

A7.42 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.

A7.43 We also asked those who switched 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.

Figure A7.8: Switching experience in the Context Consulting market research

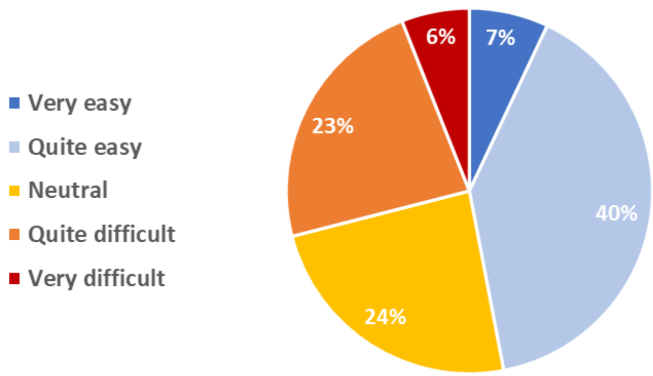

Source: Context Consulting market research report, slide 109.

[129] [✂] response dated [✂] to our cloud services questionnaire, questions [✂].
[130] [✂] 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 [✂].
[131] [✂] 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 [✂].

A7.44 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 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.[132]

A7.45 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 A7.9: Potential reasons for switching providers

Source: Context Consulting market research report, slide 110.

A7.46 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.

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

Figure A7.10: Potential reasons for adding an IaaS/PaaS provider.

Source: Context Consulting market research report, slide 111.

Challenges and obstacles to switching

A7.47 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 A7.11: Perceived challenges of switching in the Context Consulting market research

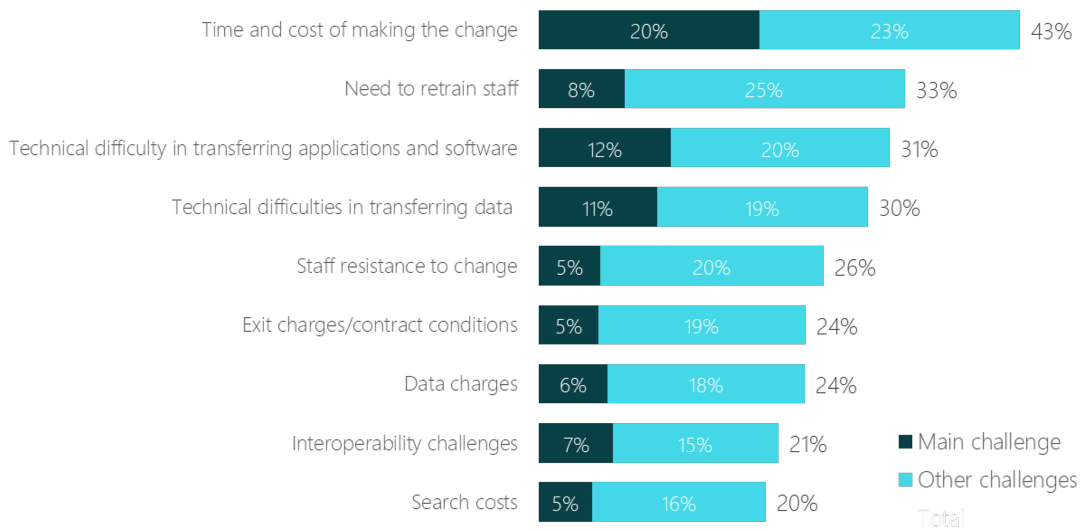

Source: Context Consulting market research report, slide 123.

A7.48 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.

A7.49 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.

A7.50 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 need to retrain staff.[133] 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].[134]

[133] [✂] response dated [✂] to the s.174 notice dated [✂], question [✂]; [✂] response dated [✂] to our cloud services questionnaire, question [✂].
[134] [✂] response dated [✂] to the s.174 notice dated [✂] questions [✂].

A8. 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.

Term	Definition
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.
De facto standard	A standard that is adopted in practice by the industry but has not undergone any formal process to obtain consensus.
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

Term	Definition
	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.
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.
Industrial internet of things (IIoT)	IoT used to enhance industrial and manufacturing business processes and applications.
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.
Internet of things (IoT)	The network of devices that contain the hardware (including sensors and actuators), software and firmware which allow

Term	Definition
	the devices to connect, interact, and freely exchange data and information.
Interoperability	The ability of computer systems or software to communicate with one another.
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.
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.
On-demand computing	On-demand computing is a delivery model in which computing resources are made available to the customer as needed, on a pay-per-use and self-service basis. The resources can be scaled up and down to address the needs of a particular workload.
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.

Term	Definition
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.
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.
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

Term	Definition
	and storage) to be divided into multiple virtual servers, commonly called virtual machines.
Virtualised network functions (VNFs)	Telco network functions (e.g. switching, routing, policy control) that are decoupled from the underlying hardware and converted to software using virtualisation.
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.